



Convergence of a Particle-based Approximation of the Block Online Expectation Maximization Algorithm

Sylvain Le Corff, Gersende Fort

► To cite this version:

Sylvain Le Corff, Gersende Fort. Convergence of a Particle-based Approximation of the Block Online Expectation Maximization Algorithm. ACM Transactions on Modeling and Computer Simulation, 2013, 23 (1). hal-00638388v3

HAL Id: hal-00638388

<https://hal.science/hal-00638388v3>

Submitted on 29 May 2012

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Convergence of a Particle-based Approximation of the Block Online Expectation Maximization Algorithm

Sylvain Le Corff^{*†} and Gersende Fort[†]

May 29, 2012

Abstract

Online variants of the Expectation Maximization (EM) algorithm have recently been proposed to perform parameter inference with large data sets or data streams, in independent latent models and in hidden Markov models. Nevertheless, the convergence properties of these algorithms remain an open problem at least in the hidden Markov case. This contribution deals with a new online EM algorithm which updates the parameter at some deterministic times. Some convergence results have been derived even in general latent models such as hidden Markov models. These properties rely on the assumption that some intermediate quantities are available in closed form or can be approximated by Monte Carlo methods when the Monte Carlo error vanishes rapidly enough. In this paper, we propose an algorithm which approximates these quantities using Sequential Monte Carlo methods. The convergence of this algorithm and of an averaged version is established and their performance is illustrated through Monte Carlo experiments.

This extended version of the paper “Convergence of a Particle-based Approximation of the Block Online Expectation Maximization Algorithm“, by S. Le Corff and G. Fort, provides detailed proofs which have been omitted in the submitted paper since they are very close to existing results. These additional proofs are postponed to Appendix B.

^{*}This work is partially supported by the French National Research Agency, under the programs ANR-07-ROBO-0002 and ANR-08-BLAN-0218.

[†]LTCI, CNRS and TELECOM ParisTech, 46 rue Barrault 75634 Paris Cedex 13, France

1 Introduction

The Expectation Maximization (EM) algorithm is a well-known iterative algorithm to solve maximum likelihood estimation in incomplete data models, see [14]. Each iteration is decomposed into two steps: in the E-step the conditional expectation of the complete log-likelihood (log of the joint distribution of the hidden states and the observations) given the observations is computed; and the M-step updates the parameter estimate. The EM algorithm is mostly practicable if the model belongs to the curved exponential family, see [29, Section 1.5] and [6, Section 10.1], so that we assume below that our model belongs to this family. Under mild regularity conditions, this algorithm is known to converge to the stationary points of the log-likelihood of the observations, see [36]. However, the original EM algorithm cannot be used to perform online estimation or when the inference task relies on large data sets. Each iteration requires the whole data set and each piece of data needs to be stored and scanned to produce a new parameter estimate. *Online* variants of the EM algorithm were first proposed for independent and identically distributed (i.i.d.) observations: [5] proposed to replace the original E-step by a stochastic approximation using the new observation. Solutions have also been proposed in hidden Markov models (HMM): [4] provides an algorithm for finite state-space HMM which relies on recursive computations of the filtering distributions combined with a stochastic approximation step. Note that, since the state-space is finite, deterministic approximations of these distributions are available. This algorithm has been extended to the case of general state-space models, the approximations of the filtering distributions being handled with Sequential Monte Carlo (SMC) algorithms, see [3], [10] and [25]. Unfortunately, it is quite challenging to address the asymptotic behavior of these algorithms (in the HMM case) since the recursive computation of the filtering distributions relies on approximations which are really difficult to control.

In [23], another online variant of the EM algorithm in HMM is proposed, called the Block Online EM (BOEM) algorithm. In this case, the data stream is decomposed into blocks of increasing sizes. Within each block, the parameter estimate is kept fixed and the update occurs at the end of the block. This update is based on a single scan of the observations, so that it is not required to store any block of observations. [23] provides results on the convergence and on the convergence rates of the BOEM algorithms. These analyses are established when the E-step (computed on each block) is available in closed form and when it can be approximated using Monte Carlo methods, under an assumption on the L_p -error of the Monte Carlo approximation.

In this paper, we consider the case when the E-step of the BOEM algorithm is computed with SMC approximations: the filtering distributions are approximated using a set of random weighted particles, see [6] and [8]. The Monte Carlo approximation is based on an online variant of the Forward Filtering Backward Smoothing algorithm (FFBS) proposed in [4] and [10]. This method is appealing for two reasons: first, it can be implemented forwards in time i.e. within a block, each observation is scanned once and never stored and the approximation

computed on each block does not require a backward step - this is crucial in our online estimation framework. Secondly, recent work on SMC approximations provides L_p -mean control of the Monte Carlo error, see e.g. [19] and [9]. This control, combined with the results in [23], sparks off the convergence results and the convergence rates provided in this contribution.

The paper is organized as follows: our new algorithm called the *Particle Block Online EM* algorithm (P-BOEM) is derived in Section 2 together with an *averaged* version. Section 3 is devoted to practical applications: the P-BOEM algorithm is used to perform parameter inference in stochastic volatility models and in the more challenging framework of the Simultaneous Localization And Mapping problem (SLAM). The convergence properties and the convergence rates of the P-BOEM algorithms are given in Section 4.

2 The Particle Block Online EM algorithms

In Section 2.1, we fix notation that will be used throughout this paper. We then derive our online algorithms in Sections 2.2 and 2.3. We finally detail, in Section 2.4, the SMC procedure that makes our algorithm a true online algorithm.

2.1 Notations and Model assumptions

A hidden Markov model on $\mathbb{X} \times \mathbb{Y}$ is defined by an initial distribution χ on $(\mathbb{X}, \mathcal{X})$ and two families of transition kernels. In this paper, the transition kernels are parametrized by $\theta \in \Theta$, where $\Theta \subseteq \mathbb{R}^{d_\theta}$ is a compact set. In the sequel, the initial distribution χ on $(\mathbb{X}, \mathcal{X})$ is assumed to be known and fixed. The parameter is estimated online in the maximum likelihood sense using a sequence of observations $\mathbf{Y}$. Online maximum likelihood parameter inference algorithms were proposed either with a gradient approach or an EM approach. In the case of finite state-spaces HMM, [26] proposed a recursive maximum likelihood procedure. The asymptotic properties of this algorithm have recently been addressed in [34]. This algorithm has been adapted to general state-spaces HMM with SMC methods (see [18]). The main drawback of gradient methods is the necessity to scale the gradient components. As an alternative to performing online inference in HMM, online EM based algorithms have been proposed for finite state-spaces (see [4]) or general state-spaces HMM (see [3], [10] and [23]). [10] proposed a SMC method giving encouraging experimental results. Nevertheless, it relies on a combination of stochastic approximations and SMC computations so that its analysis is quite challenging. In [23], the convergence of an online EM based algorithm is established. This algorithm requires either the exact computation of intermediate quantities (available explicitly only in finite state-spaces HMM or in linear Gaussian models) or the use of Monte Carlo methods to approximate these quantities. We propose to apply this algorithm to general models where these quantities are replaced by SMC approximations. We prove that the Monte Carlo error is controlled in such a way that the convergence

properties of [23] hold for the P-BOEM algorithms.

We now detail the model assumptions. Consider a family of transition kernels $\{m_\theta(x, x')d\lambda(x')\}_{\theta \in \Theta}$ on $\mathbb{X} \times \mathcal{X}$, where $\mathbb{X}$ is a general state-space equipped with a countably generated σ -field $\mathcal{X}$, and λ is a finite measure on $(\mathbb{X}, \mathcal{X})$. Let $\{g_\theta(x, y)d\nu(y)\}_{\theta \in \Theta}$ be a family of transition kernels on $\mathbb{X} \times \mathcal{Y}$, where $\mathbb{Y}$ is a general space endowed with a countably generated σ -field $\mathcal{Y}$ and ν is a measure on $(\mathbb{Y}, \mathcal{Y})$. Let $\mathbf{Y} = \{\mathbf{Y}_t\}_{t \in \mathbb{Z}}$ be the observation process defined on $(\Omega, \mathbb{P}, \mathcal{F})$ and taking values in $\mathbb{Y}^{\mathbb{Z}}$. The batch EM algorithm is an offline maximum likelihood procedure which iteratively produces parameter estimates using the complete data log-likelihood (log of the joint distribution of the observations and the states) and a fixed set of observations, see [14]. In the HMM context presented above, given T observations $\mathbf{Y}_{1:T}$, the missing data $x_{0:T}$ and a parameter θ , the complete data log-likelihood may be written as (up to the initial distribution χ which is assumed to be known)

$$\ell_\theta(x_{0:T}, \mathbf{Y}_{1:T}) \stackrel{\text{def}}{=} \sum_{t=1}^T \{\log m_\theta(x_{t-1}, x_t) + \log g_\theta(x_t, \mathbf{Y}_t)\} , \quad (1)$$

where we use $x_{r:t}$ as a shorthand notation for the sequence $(x_r, \dots, x_t)$, $r \leq t$. Each iteration of the batch EM algorithm is decomposed into two steps. The E-step computes, for all $\theta \in \Theta$, an expectation of the complete data log-likelihood under the conditional probability of the hidden states given the observations and the current parameter estimate $\hat{\theta}$. In the HMM context, due to the additive form of the complete data log-likelihood (1), the E-step is decomposed into T expectations under the conditional probabilities $\Phi_{\hat{\theta}, t, T}^{\chi, 0}(\cdot, \mathbf{y})$ where

$$\Phi_{\hat{\theta}, s, t}^{\chi, r}(h, \mathbf{y}) \stackrel{\text{def}}{=} \frac{\int \chi(dx_r) \{\prod_{i=r}^{t-1} m_{\hat{\theta}}(x_i, x_{i+1}) g_{\hat{\theta}}(x_{i+1}, \mathbf{y}_{i+1})\} h(x_{s-1}, x_s, \mathbf{y}_s) d\lambda(x_{r+1:t})}{\int \chi(dx_r) \{\prod_{i=r}^{t-1} m_{\hat{\theta}}(x_i, x_{i+1}) g_{\hat{\theta}}(x_{i+1}, \mathbf{y}_{i+1})\} d\lambda(x_{r+1:t})} , \quad (2)$$

for any bounded function h , any $\theta \in \Theta$, any $r < s \leq t$ and any sequence $\mathbf{y} \in \mathbb{Y}^{\mathbb{Z}}$. Then, given the current value of the parameter $\hat{\theta}$, the E-step amounts to computing the quantity

$$Q_T(\theta, \hat{\theta}) \stackrel{\text{def}}{=} \frac{1}{T} \sum_{t=1}^T \Phi_{\hat{\theta}, t, T}^{\chi, 0}(\log m_\theta + \log g_\theta, \mathbf{Y}) , \quad (3)$$

for any $\theta \in \Theta$. The M-step sets the new parameter estimate as a maximum of this expectation over θ .

The computation of $\theta \mapsto Q_T(\theta, \hat{\theta})$ for any $\theta \in \Theta$ is usually intractable except in the case of complete data likelihood belonging to the curved exponential family, see [29, Section 1.5] and [6, Section 10.1]. Therefore, in the sequel, the following assumption is assumed to hold:

- A1** (a) There exist continuous functions $\phi : \Theta \rightarrow \mathbb{R}$, $\psi : \Theta \rightarrow \mathbb{R}^d$ and $S : \mathbb{X} \times \mathbb{X} \times \mathbb{Y} \rightarrow \mathbb{R}^d$ s.t.

$$\log m_\theta(x, x') + \log g_\theta(x', y) = \phi(\theta) + \langle S(x, x', y), \psi(\theta) \rangle ,$$

where $\langle \cdot, \cdot \rangle$ denotes the scalar product on $\mathbb{R}^d$.

- (b) There exists an open subset $\mathcal{S}$ of $\mathbb{R}^d$ that contains the convex hull of $S(\mathbb{X} \times \mathbb{X} \times \mathbb{Y})$.
- (c) There exists a continuous function $\bar{\theta} : \mathcal{S} \rightarrow \Theta$ s.t. for any $s \in \mathcal{S}$,

$$\bar{\theta}(s) = \operatorname{argmax}_{\theta \in \Theta} \{ \phi(\theta) + \langle s, \psi(\theta) \rangle \} .$$

Under A1, the quantity $Q_T(\theta, \hat{\theta})$ defined by (3) becomes

$$Q_T(\theta, \hat{\theta}) = \phi(\theta) + \left\langle \frac{1}{T} \sum_{t=1}^T \Phi_{\hat{\theta}, t, T}^{\chi, 0}(S, \mathbf{Y}), \psi(\theta) \right\rangle , \quad (4)$$

so that the definition of the function $\theta \mapsto Q_T(\theta, \hat{\theta})$ requires the computation of an expectation $\frac{1}{T} \sum_{t=1}^T \Phi_{\hat{\theta}, t, T}^{\chi, 0}(S, \mathbf{Y})$ independently of θ .

The M-step of the batch EM iteration amounts to computing

$$\bar{\theta} \left(\frac{1}{T} \sum_{t=1}^T \Phi_{\hat{\theta}, t, T}^{\chi, 0}(S, \mathbf{Y}) \right) .$$

This batch EM algorithm is designed for a fixed set of observations. A natural extension of this algorithm to the online context is to define a sequence of parameter estimates by

$$\theta_{t+1} = \operatorname{argmax}_{\theta} Q_{t+1}(\theta, \theta_t) .$$

Unfortunately, the computation of $Q_{t+1}(\theta, \theta_t)$ requires the whole set of observations to be stored and scanned for each estimation. For large data sets the computation cost of the E-step makes it intractable in this case. To overcome this difficulty, several online variants of the batch EM algorithm have been proposed, based on a recursive approximation of the function $\theta \mapsto Q_{t+1}(\cdot, \theta_t)$ (see [3], [10] and [23]). In this paper, we focus on the Block Online EM (BOEM) algorithm, see [23].

2.2 Particle Block Online EM (P-BOEM)

The BOEM algorithm, introduced in [23], is an online variant of the EM algorithm. The observations are processed sequentially per block and the parameter estimate is updated at the end of each block. Let $\{\tau_k\}_{k \geq 1}$ be a sequence of positive integers denoting the length of the blocks and set

$$T_n \stackrel{\text{def}}{=} \sum_{k=1}^n \tau_k \quad \text{and} \quad T_0 \stackrel{\text{def}}{=} 0 ; \quad (5)$$

$\{T_k\}_{k \geq 1}$ are the deterministic times at which the parameter updates occur. Define, for all integers $\tau > 0$ and $T \geq 0$ and all $\theta \in \Theta$,

$$\bar{S}_\tau^{\chi, T}(\theta, \mathbf{Y}) \stackrel{\text{def}}{=} \frac{1}{\tau} \sum_{t=T+1}^{T+\tau} \Phi_{\theta, t, T+\tau}^{\chi, T}(S, \mathbf{Y}) . \quad (6)$$

The quantity $\bar{S}_\tau^{\chi, T}(\theta, \mathbf{Y})$ corresponds to the intermediate quantity in (4) with the observations $\mathbf{Y}_{T+1:T+\tau}$.

The BOEM algorithm iteratively defines a sequence of parameter estimates $\{\theta_n\}_{n \geq 0}$ as follows: given the current parameter estimate θ_n ,

- (i) compute the quantity $\bar{S}_{\tau_{n+1}}^{\chi, T_n}(\theta_n, \mathbf{Y})$,
- (ii) compute a candidate for the new value of the parameter: $\theta_{n+1} = \bar{\theta} \left(\bar{S}_{\tau_{n+1}}^{\chi, T_n}(\theta_n, \mathbf{Y}) \right)$,

To make the exposition easier, we assume that the initial distribution χ is the same on each block. The dependence of $\bar{S}_\tau^{\chi, T}(\theta, \mathbf{Y})$ on χ is thus dropped from the notation for better clarity.

The quantity $\bar{S}_{\tau_{n+1}}^{T_n}(\theta_n, \mathbf{Y})$ is available in closed form only in the case of linear Gaussian models and HMM with finite state-spaces. In HMM with general state-spaces $\bar{S}_{\tau_{n+1}}^{T_n}(\theta_n, \mathbf{Y})$ cannot be computed explicitly and we propose to compute an approximation of $\bar{S}_{\tau_{n+1}}^{T_n}(\theta_n, \mathbf{Y})$ using SMC algorithms thus yielding the *Particle-BOEM* (P-BOEM) algorithm. Different methods can be used to compute these approximations (see e.g. [9], [10] and [16]). We will discuss in Section 2.4 below some SMC approximations that use the data sequentially.

Denote by $\tilde{S}_n(\theta, \mathbf{Y})$ the SMC approximation of $\bar{S}_{\tau_{n+1}}^{T_n}(\theta, \mathbf{Y})$ computed with N_{n+1} particles. The P-BOEM algorithm iteratively defines a sequence of parameter estimates $\{\theta_n\}_{n \geq 0}$ as follows: given the current parameter estimate θ_n ,

- (i) compute the quantity $\tilde{S}_n(\theta_n, \mathbf{Y})$,
- (ii) compute a candidate for the new value of the parameter:

$$\theta_{n+1} = \bar{\theta} \left(\tilde{S}_n(\theta_n, \mathbf{Y}) \right) .$$

We give in Algorithm 1 lines 1 to 9 an algorithmic description of the P-BOEM algorithm. Note that the idea of processing the observations by blocks is proposed in [30] to fit a normal mixture model. The incremental EM algorithm discussed in [30] is an alternative to the batch EM algorithm for very large data sets. Contrary to our framework, in the algorithm proposed by [30], the number of observations is fixed and the same observations are scanned several times.

2.3 Averaged Particle Block Online EM

Following the same lines as in [23], we propose to replace the P-BOEM sequence $\{\theta_n\}_{n \geq 0}$ by an *averaged* sequence. This new sequence can be computed recursively, simultaneously with the P-BOEM sequence, and does not require additional storage of the data. The proposed averaged P-BOEM algorithm is defined as follows (see also lines 5 and 6 of Algorithm 1): the step (ii) of the P-BOEM algorithm presented above is followed by

(iv) compute the quantity

$$\tilde{\Sigma}_{n+1} = \frac{T_n}{T_{n+1}} \tilde{\Sigma}_n + \frac{\tau_{n+1}}{T_{n+1}} \tilde{S}_n(\theta_n, \mathbf{Y}) , \quad (7)$$

(v) define

$$\tilde{\theta}_{n+1} \stackrel{\text{def}}{=} \bar{\theta} \left(\tilde{\Sigma}_{n+1} \right) . \quad (8)$$

We set $\tilde{\Sigma}_0 = 0$ so that

$$\tilde{\Sigma}_n = \frac{1}{T_n} \sum_{j=1}^n \tau_j \tilde{S}_{j-1}(\theta_{j-1}, \mathbf{Y}) ; \quad (9)$$

we will prove in Section 4.4 that the rate of convergence of the averaged sequence $\{\tilde{\theta}_n\}_{n \geq 0}$, computed from the averaged statistics $\{\tilde{\Sigma}_n\}_{n \geq 0}$, is better than the non-averaged one. We will also observe this property in Section 3 by comparing the variability of the P-BOEM and the averaged P-BOEM sequences in numerical applications.

Algorithm 1 P-BOEM and averaged P-BOEM

Require: $\theta_0, \{\tau_n\}_{n \geq 1}, \{N_n\}_{n \geq 1}, \{\mathbf{Y}_t\}_{t \geq 0}$.

Ensure: $\{\theta_n\}_{n \geq 0}$ and $\{\tilde{\theta}_n\}_{n \geq 0}$.

Set $\tilde{\Sigma}_0 = 0$.

for all $i \geq 0$ **do**

 Compute sequentially $\tilde{S}_i(\theta_i, \mathbf{Y})$.

 Set $\theta_{i+1} = \bar{\theta} \left(\tilde{S}_i(\theta_i, \mathbf{Y}) \right)$.

 Set

$$\tilde{\Sigma}_{i+1} = \frac{T_i}{T_{i+1}} \tilde{\Sigma}_i + \frac{\tau_{i+1}}{T_{i+1}} \tilde{S}_i(\theta_i, \mathbf{Y}) .$$

 Set $\tilde{\theta}_{i+1} = \bar{\theta} \left(\tilde{\Sigma}_{i+1} \right)$.

end for

2.4 The SMC approximation step

As the P-BOEM algorithm is an online algorithm, the SMC algorithm should use the data sequentially: no backward pass is allowed to browse all the data at the end of the block. Hence, the approximation is computed recursively within each block, each observation being used once and never stored. These SMC algorithms will be referred to as *forward only SMC*. We detail below a forward only SMC algorithm for the computation of $\tilde{S}_n(\theta_n, \mathbf{Y})$ which has been proposed by [4] (see also [10]).

For notational convenience, the dependence on n is omitted. For block n , the algorithm below has to be applied with $(\tau, N) \leftarrow (\tau_{n+1}, N_{n+1})$, $Y_{1:\tau} \leftarrow Y_{T_n+1, T_n+\tau_{n+1}}$ and $\theta \leftarrow \theta_n$.

The key property is to observe that

$$\bar{S}_\tau^0(\theta, \mathbf{Y}) = \phi_\tau^\theta(R_{\theta, \tau}) \quad (10)$$

where ϕ_t^θ is the filtering distribution at time t , and the functions $R_{t, \theta} : \mathbb{X} \rightarrow \mathcal{S}$, $1 \leq t \leq \tau$, satisfy the equations

$$R_{t, \theta}(x) = \frac{1}{t} B_t^\theta(x, S(\cdot, x, Y_t)) + \frac{t-1}{t} B_t^\theta(x, R_{t-1, \theta}) , \quad (11)$$

where B_t^θ denotes the backward smoothing kernel at time t

$$B_t^\theta(x, dx') = \frac{m_\theta(x', x)}{\int m_\theta(u, x) \phi_{t-1}^\theta(du)} \phi_{t-1}^\theta(dx') . \quad (12)$$

By convention, $R_{0, \theta}(x) = 0$ and $\phi_0^\theta = \chi$. A proof of the equalities (10) to (12) can be found in [4] and [10]. Therefore, a careful reading of Eqs (10) to (12) shows that, for an iterative particle approximation of $\bar{S}_\tau^0(\theta, \mathbf{Y})$, it is sufficient to update from time $t-1$ to t

- (i) N weighted samples $\{(\xi_t^\ell, \omega_t^\ell); \ell \in \{1, \dots, N\}\}$ used to approximate the filtering distribution ϕ_t^θ .
- (ii) the intermediate quantities $\{R_{t, \theta}^\ell\}_{\ell=1}^N$, approximating the function $R_{t, \theta}$ at point $x = \xi_t^\ell$, $\ell \in \{1, \dots, N\}$.

We describe below such an algorithm. An algorithmic description is also provided in Appendix A, Algorithm 2.

Given *instrumental* Markov transition kernels $\{q_t(x, x'), t \leq \tau\}$ on $\mathbb{X} \times \mathcal{X}$ and adjustment multipliers $\{v_t, t \leq \tau\}$, the procedure goes as follows:

- (i) *line 1 in Algorithm 2*: sample independently N particles $\{\xi_0^\ell\}_{\ell=1}^N$ with the same distribution χ .
- (ii) *line 6 in Algorithm 2*: at each time step $t \in \{1, \dots, \tau\}$, pairs $\{(J_t^\ell, \xi_t^\ell)\}_{\ell=1}^N$ of indices and particles are sampled independently (conditionally to $Y_{1:t}$, θ and $\{(J_{t-1}^\ell, \xi_{t-1}^\ell)\}_{\ell=1}^N$) from the instrumental distribution:

$$\pi_t(i, dx) \propto \omega_{t-1}^i v_t(\xi_{t-1}^i) q_t(\xi_{t-1}^i, x) \lambda(dx) , \quad (13)$$

on the product space $\{1, \dots, N\} \times \mathbb{X}$. For any $t \in \{1, \dots, \tau\}$ and any $\ell \in \{1, \dots, N\}$, J_t^ℓ denotes the index of the selected particle at time $t-1$ used to produce ξ_t^ℓ .

- (iii) *line 7 in Algorithm 2*: once the new particles $\{\xi_t^\ell\}_{\ell=1}^N$ have been sampled, their importance weights $\{\omega_t^\ell\}_{\ell=1}^N$ are computed.
- (iv) *lines 8 in Algorithm 2*: update the intermediate quantities $\{R_{t, \theta}^\ell\}_{\ell=1}^N$.

If, for all $x \in \mathbb{X}$, $v_t(x) = 1$ and if the kernels q_t are chosen such that $q_t = m_\theta$, lines 6-7 in Algorithm 2 are known as the *Bootstrap filter*. Other choices of q_t and v_t can be made, see e.g. [6].

3 Applications to Bayesian inverse problems in Hidden Markov Models

3.1 Stochastic volatility model

Consider the following stochastic volatility model:

$$X_{t+1} = \phi X_t + \sigma U_t, \quad Y_t = \beta e^{\frac{X_t}{2}} V_t,$$

where $X_0 \sim \mathcal{N}(0, (1 - \phi^2)^{-1} \sigma^2)$ and $\{U_t\}_{t \geq 0}$ and $\{V_t\}_{t \geq 0}$ are two sequences of i.i.d. standard Gaussian r.v., independent from X_0 .

We illustrate the convergence of the P-BOEM algorithms and discuss the choice of some design parameters such as the pair (τ_n, N_n) . Data are sampled using $\phi = 0.95$, $\sigma^2 = 0.1$ and $\beta^2 = 0.6$; we estimate $\theta = (\phi, \sigma^2, \beta^2)$ by applying the P-BOEM algorithm and its averaged version. All runs are started from $\phi = 0.1$, $\sigma^2 = 0.6$ and $\beta^2 = 2$.

Figure 1 displays the estimation of the three parameters as a function of the number of observations, over 50 independent Monte Carlo runs. The block-size sequence is of the form $\tau_n \propto n^{1.2}$. For the SMC step, we choose $N_n = 0.25 \cdot \tau_n$; particles are sampled as described in Algorithm 2 (see Appendix A) with the bootstrap filter. For each parameter, Figure 1 displays the empirical median (bold line) and upper and lower quartiles (dotted line). The averaging procedure is started after 1500 observations. Both algorithms converge to the true values of the parameters and, once the averaging procedure is started, the variance of the estimation decreases (estimation of ϕ and β^2). The estimation of σ^2 shows that, if the averaging procedure is started with too few observations, the estimation can be slowed down.

We now discuss the role of the pairs (τ_n, N_n) . Roughly speaking (see section 4 for a rigorous decomposition), τ controls the rate of convergence of $\bar{S}_\tau^T(\theta, \mathbf{Y})$ to $\lim_{\tau \rightarrow \infty} \bar{S}_\tau^T(\theta, \mathbf{Y})$; and N controls the error between $\bar{S}_\tau^T(\theta, \mathbf{Y})$ and its SMC approximation. We will show in Section 4 that $\lim_n \tau_n = \lim_n N_n = +\infty$ are part of some sufficient conditions for the P-BOEM algorithms to converge. We thus choose increasing sequences $\{\tau_n, N_n\}_{n \geq 1}$. The role of τ_n has been illustrated in [23, Section 3]. Hence, in this illustration, we fix τ_n and discuss the role of N_n . Figure 2 compares the algorithms when applied with $\tau_n \propto n^{1.1}$ and $N_n = \sqrt{\tau_n}$ or $N_n = \tau_n$. The empirical variance (over 50 independent Monte Carlo runs) of the estimation of β^2 is displayed, as a function of the number of blocks. First, Figure 2 illustrates the variance decrease provided by the averaged procedure, whatever the block size sequence. Moreover, increasing the number of particles per block improves the variance of the estimation given by the P-BOEM algorithm while the impact on the variance of the averaged estimation is less important. On average, the variance is reduced by a factor of 3.0 for the P-BOEM algorithm and by a factor of 1.8 for its averaged version when the number of particles goes from $N_n = \sqrt{\tau_n}$ to $N_n = \tau_n$. These practical considerations illustrate the theoretical results derived in Section 4.4.

Finally, we discuss the role of the initial distribution χ . In all the applications






Figure 1: Estimation of ϕ , σ^2 and β^2 without (left) and with (right) averaging. Each graph represents the empirical median (bold line) and upper and lower quartiles (dotted line) over 50 independent Monte Carlo runs. The averaging procedure is started after 1500 observations. The first 1000 observations are not displayed for better clarity.

above, we have the same distribution $\chi \equiv \mathcal{N}(0, (1 - \phi^2)^{-1} \sigma^2)$ at the beginning of each block. We could choose a different distribution χ_n for each block such as, e.g., the filtering distribution at the end of the previous block. We have observed that this particular choice of χ_n leads to the same behavior for both algorithms.

To end this section, the P-BOEM algorithm is compared to the Online EM algorithm outlined in [4] and [10]. These algorithms rely on a combination of stochastic approximation and SMC methods. According to classical results on stochastic approximation, it is expected that the rate of convergence of the



(a) **P-BOEM**: empirical variance of the estimation of β^2 with $N_n = \sqrt{\tau_n}$ (dashed line) and $N_n = \tau_n$ (bold line).



(b) **Averaged P-BOEM**: empirical variance of the estimation of β^2 with $N_n = \sqrt{\tau_n}$ (dashed line) and $N_n = \tau_n$ (bold line).

Figure 2: Empirical variance of the estimation of β^2 with the P-BOEM algorithm (top) and its averaged version (bottom). The averaging procedure is started after the 25-th block and the variance is displayed after a burn-in time of 35 blocks.

Online EM algorithm behaves like $\gamma_n^{1/2}$, where $\{\gamma_n\}_{n \geq 0}$ is the so called step-size sequence. Hence, γ_n in the Online EM algorithm is chosen such that $\gamma_n \propto n^{-0.55}$ and the block-size sequence in the P-BOEM algorithm such that $\tau_n \propto n^{1.2}$. The number of particles used in the Online EM algorithm is fixed and chosen so that the computational costs of both algorithms are similar. Provided that $N_n \propto \tau_n$ in the P-BOEM algorithm, this leads to a choice of 70 particles for the Online EM algorithm. We report in Figure 3, the estimation of ϕ and σ^2 for a Polyak-Ruppert averaged Online EM algorithm (see [33]) and the averaged P-BOEM algorithm as a function of the number of observations. The averaging procedure is started after about 1500 observations. As noted in [23, Section 3] for a constant sequence $\{N_n\}_{n \geq 0}$ this figure shows that both algorithms behave

similarly. For the estimation of ϕ and β^2 , the variance is smaller for the P-BOEM algorithm and the convergence is faster for the P-BOEM algorithm in the case of β^2 . Conclusions are different for the estimation of σ^2 : the variance is smaller for the P-BOEM algorithm but the Online EM algorithm converges a bit faster. The main advantage of the P-BOEM algorithm is that it relies on approximations which can be controlled in such a way that we are able to show that the limiting points of the P-BOEM algorithms are the stationary points of the limiting normalized log-likelihood of the observations.

3.2 Simultaneous Localization And Mapping

The Simultaneous Localization And Mapping (SLAM) problem arises when a mobile device wants to build a map of an unknown environment and, at the same time, has to estimate its position in this map. The common statistical approach for the SLAM problem is to introduce a state-space model. Many solutions have been proposed depending on the assumptions made on the transition and observation models, and on the map (see e.g. [2], [28] and [32]). In [28] and [25], it is proposed to see the SLAM as an inference problem in HMM: the localization of the robot is the hidden state with Markovian dynamic, and the map is seen as an unknown parameter. Therefore, the mapping problem is answered by solving the inference task, and the localization problem is answered by approximating the conditional distribution of the hidden states given the observations.

In this application, we consider a statistical model for a landmark-based SLAM problem for a bicycle manoeuvring on a plane surface.

Let $x_t \stackrel{\text{def}}{=} \{x_{t,i}\}_{i=1}^3$ be the robot position, where $x_{t,1}$ and $x_{t,2}$ are the robot's cartesian coordinates and $x_{t,3}$ its orientation. At each time step, deterministic controls are sent to the robot so that it explores a given part of the environment. Controls are denoted by (v_t, ψ_t) where ψ_t stands for the robot's heading direction and v_t its velocity. The robot position at time t , given its previous position at time $t-1$ and the noisy controls $(\hat{v}_t, \hat{\psi}_t)$, can be written as


$$x_t = f(x_{t-1}, \hat{v}_t, \hat{\psi}_t), \quad (14)$$

where $(\hat{v}_t, \hat{\psi}_t)$ is a 2-dimensional Gaussian distribution with mean (v_t, ψ_t) and known covariance matrix Q . In this contribution we use the kinematic model of the front wheel of a bicycle (see e.g. [1]) where the function f in (14) is given by


$$f(x_{t-1}, \hat{v}_t, \hat{\psi}_t) = x_{t-1} + \begin{pmatrix} \hat{v}_t d_t \cos(x_{t-1,3} + \hat{\psi}_t) \\ \hat{v}_t d_t \sin(x_{t-1,3} + \hat{\psi}_t) \\ \hat{v}_t d_t B^{-1} \sin(\hat{\psi}_t) \end{pmatrix},$$

where d_t is the time period between two successive positions and B is the robot wheelbase.


The 2-dimensional environment is represented by a set of landmarks $\theta \stackrel{\text{def}}{=} \{\theta_j\}_{1 \leq j \leq q}$, $\theta_j \in \mathbb{C}$ being the position of the j -th landmark. The total number of landmarks q and the association between observations and landmarks are assumed to be known.



(a) Estimation of ϕ .



(b) Estimation of σ^2 .



(c) Estimation of β^2 .

Figure 3: Estimation of ϕ , σ^2 and β^2 with the averaged P-BOEM algorithm (left) and a Polyak-Ruppert averaged version of the Online EM algorithm (right) after 300, 1500, 5000, 10000, 20000 and 45000 observations. The averaging procedure is started after about 1000 observations (which corresponds to the 25-th block for the P-BOEM algorithm).

At time t , the robot observes the distance and the angular position of all landmarks in its neighborhood; let $c_t \subseteq \{1, \dots, q\}$ be the set of observed landmarks at time t . It is assumed that the observations $\{y_{t,i}\}_{i \in c_t}$ are independent

and satisfy

$$y_{t,i} = h(x_t, \theta_i) + \delta_{t,i} ,$$

where h is defined by

$$h(x, \kappa) \stackrel{\text{def}}{=} \begin{pmatrix} \sqrt{(\kappa_1 - x_1)^2 + (\kappa_2 - x_2)^2} \\ \arctan \frac{\kappa_2 - x_2}{\kappa_1 - x_1} - x_3 \end{pmatrix} ,$$

and the noise vectors $\{\delta_{t,i}\}_{t,i}$ are i.i.d Gaussian $\mathcal{N}(0, R)$. R is assumed to be known.

The model presented in this Section does not take into account all the issues arising in the SLAM problem (such as the association process which is assumed to be known and the known covariance matrices). The aim is to prove that the BOEM algorithm and its averaged version have satisfying behavior even in the challenging framework described above. The observation and motion models are highly nonlinear and we show that the BOEM algorithm remains stable in this experiment. Several solutions have been proposed to solve the association problem (see e.g. [2] for a solution based on the likelihood of the observations) and could be adapted to our case. We want to estimate $\theta = \{\theta_j\}_{j=1}^q$ by applying the P-BOEM algorithms. In this paper, we use simulated data. $q = 15$ landmarks are drawn in a square of size $45m \times 45m$. The robot path is sampled with a given set of controls. Using the *true* positions of all landmarks in the map and the true path of the robot (see the dots and the bold line on Figure 4), observations are sampled by setting: $R = \begin{pmatrix} \sigma_r^2 & \rho \\ \rho & \sigma_b^2 \end{pmatrix}$, where $\sigma_r = 0.5m$, $\sigma_b = \frac{\pi}{60} \text{rad}$ and $\rho = 0.01$. We choose $Q = \text{diag}(\sigma_v^2, \sigma_\phi^2)$ where $\sigma_v = 0.5m.s^{-1}$, $\sigma_\phi = \frac{\pi}{60} \text{rad}$ and $B = 1.5m$.

In this model, the transition denoted by m_θ does not depend on the map θ (see (14)) and the marginal likelihood g_θ is such that the complete data likelihood does not belong to the curved exponential family:

$$\sum_{i \in c_t} \ln g_\theta(x_t, y_{t,i}) \propto \sum_{i \in c_t} [y_{t,i} - h(x_t, \theta_i)]^* R^{-1} [y_{t,i} - h(x_t, \theta_i)] . \quad (15)$$

Hence, in order to apply Algorithm 1, at the beginning of each block, g_θ is approximated by a function depending on the current parameter estimate so that the resulting approximated model belongs to the curved exponential family (see [25]). As can be seen from (15), approximating the function $\kappa \mapsto h(x, \kappa)$ by its first-order Taylor expansion at θ_i leads to a quadratic approximation of g_θ . This approach is commonly used in the SLAM framework to use the properties of linear Gaussian models (see e.g. [2]).

As the landmarks are not observed all the time, we choose a slowly increasing sequence $\{\tau_n \propto n^{1.1}\}_{n \geq 1}$ so that the number of updates is not too small (in this experiment, we have 60 updates for a total number of observations of 2000). As the total number of observations is not so large (the largest block is of length 60), the number of particles is chosen to be constant on each block: for all $n \geq 1$, $N_n = 50$. For the SMC step, we apply Algorithm 2 with the *bootstrap filter*.

For each run the estimated path (equal to the weighted mean of the particles) and the estimated map at the end of the loop ($T = 2000$) are stored. Figure 4 represents the mean estimated path and the mean map over 50 independent Monte Carlo runs. It highlights the good performance of the P-BOEM algorithm in a more complex framework.




Figure 4: True trajectory (bold line) and true landmark positions (balls) with the estimated path (dotted line) and the landmarks' estimated positions (stars) at the end of the run ($T = 2000$).

We also compare our algorithm to the *marginal SLAM* algorithm proposed by [28]. In this algorithm, the map is also modeled as a parameter to learn in a HMM model; SMC methods are used to estimate the map in the maximum likelihood sense. The Marginal SLAM algorithm is a gradient-based approach for solving the recursive maximum likelihood procedure. Note that, in the case of i.i.d. observations, [35] proposed to update the parameter estimate each time a new observation is available using a stochastic gradient approach. Figure 5 illustrates the estimation of the position of each landmark. The P-BOEM algorithm is applied using the same parameters as above and the *marginal SLAM* algorithm uses a sequence of step-size $\{\gamma_n \propto n^{-0.6}\}_{n \geq 1}$. We use the averaged version of the P-BOEM algorithm and a Polyak-Ruppert based averaging procedure for the *marginal SLAM* algorithm (see [33]). For each landmark the last estimation (at the end of the loop) of the position is stored for each of the 50 independent Monte Carlo runs. Figure 5 displays the distance between the estimated position and the true position for each landmark. In this experiment, the P-BOEM based SLAM algorithm outperforms the *marginal SLAM* algorithm.

4 Convergence of the Particle Block Online EM algorithms

In this section, we analyze the limiting points of the P-BOEM algorithm. We prove in Theorem 4.3 that the P-BOEM algorithm has the same limit points as a so-called *limiting EM* algorithm, which defines a sequence $\{\theta_n\}_{n \geq 0}$ by $\theta_{n+1} = \bar{\theta}[\bar{S}(\theta_n)]$ where $\bar{S}(\theta)$ is the a.s. limit $\lim_{\tau \rightarrow +\infty} \bar{S}_\tau^T(\theta, \mathbf{Y})$ (defined by (6)). As discussed in [23, Section 4.3.], the set of limit points of the limiting EM




Figure 5: Distance between the final estimation and the true position for each of the 15 landmarks with the averaged *marginal SLAM* algorithm (left) and the averaged P-BOEM algorithm (right).

algorithm is the set of stationary points of the contrast function $\ell(\theta)$, defined as the a.s. limit of the normalized log-likelihood of the observations, when $T \rightarrow +\infty$. The convergence result below on the P-BOEM algorithm requires two sets of assumptions: conditions A2 to A5 are the same as in [23] and imply the convergence of the BOEM algorithm; assumptions A6 and A7 are introduced to control the Monte Carlo error.

4.1 Assumptions

Consider the following assumptions

A2 There exist σ_- and σ_+ s.t. for any $(x, x') \in \mathbb{X}^2$ and any $\theta \in \Theta$, $0 < \sigma_- \leq m_\theta(x, x') \leq \sigma_+$. Set $\rho \stackrel{\text{def}}{=} 1 - (\sigma_-/\sigma_+)$.

Define, for all $y \in \mathbb{Y}$,

$$b_-(y) \stackrel{\text{def}}{=} \inf_{\theta \in \Theta} \int g_\theta(x, y) \lambda(dx) \quad \text{and} \quad b_+(y) \stackrel{\text{def}}{=} \sup_{\theta \in \Theta} \int g_\theta(x, y) \lambda(dx). \quad (16)$$

For any sequence of r.v. $Z \stackrel{\text{def}}{=} \{Z_t\}_{t \in \mathbb{Z}}$ on $(\Omega, \tilde{\mathbb{P}}, \mathcal{F})$, let

$$\mathcal{F}_k^Z \stackrel{\text{def}}{=} \sigma(\{Z_u\}_{u \leq k}) \quad \text{and} \quad \mathcal{G}_k^Z \stackrel{\text{def}}{=} \sigma(\{Z_u\}_{u \geq k}) \quad (17)$$

be σ -fields associated to Z . We also define the mixing coefficients by, see [7],

$$\beta^Z(n) \stackrel{\text{def}}{=} \sup_{u \in \mathbb{Z}} \beta(\mathcal{G}_{u+n}^Z, \mathcal{F}_u^Z), \forall n \geq 0, \quad (18)$$

where for any σ -algebras $\mathcal{F}$ and $\mathcal{G}$,

$$\beta(\mathcal{G}, \mathcal{F}) \stackrel{\text{def}}{=} \sup_{B \in \mathcal{G}} |\tilde{\mathbb{P}}(B|\mathcal{F}) - \tilde{\mathbb{P}}(B)|. \quad (19)$$

For $p > 0$ and Z a $\mathbb{R}^d$ -valued random variable measurable w.r.t. the σ -algebra $\mathcal{F}$, set

$$\|Z\|_p \stackrel{\text{def}}{=} (\mathbb{E} [|Z|^p])^{1/p}.$$

$$\mathbf{A3}\text{-(}p\text{)} \left\| \sup_{x, x' \in \mathbb{X}^2} |S(x, x', \mathbf{Y}_0)| \right\|_p < +\infty.$$

A4 (a) $\mathbf{Y}$ is a stationary sequence such that there exist $C \in [0, 1]$ and $\beta \in (0, 1)$ satisfying, for any $n \geq 0$, $\beta^{\mathbf{Y}}(n) \leq C\beta^n$, where $\beta^{\mathbf{Y}}$ is defined in (18).

$$(b) \mathbb{E} [|\log b_-(\mathbf{Y}_0)| + |\log b_+(\mathbf{Y}_0)|] < +\infty.$$

A5 There exist $c > 0$ and $a > 1$ such that for all $n \geq 1$, $\tau_n = \lfloor cn^a \rfloor$.

Assumptions A2 to A5 are the same as in [23]. A2, referred to as the *strong mixing condition*, is used to prove the uniform forgetting property of the initial condition of the filter, see e.g. [11] and [12]. This assumption is easy to check in finite state-space HMM or when the state-space is compact when the Markov kernel $\mathbf{m}_\theta$ is sufficiently regular. As noted in [23], it can fail to hold in quite general situations. Nevertheless, the exponential forgetting property needed to ensure the convergence results could be checked under weaker assumptions (see [15] for a Doeblin assumption). However, it would imply quite technical supplementary results out of the scope of this paper. Examples of observation sequences satisfying A4 include, for example, stationary ψ -irreducible and positive recurrent Markov chains which are geometrically ergodic (see e.g. [31] for Markov chains theory).

We need to control the L_p -mean error on each block between $\bar{S}_{\tau_{n+1}}^{T_n}(\theta_n, \mathbf{Y})$ and its SMC approximation. This control is discussed in Section 4.2 below when the SMC approximation is computed as described in Section 2.4.

4.2 L_p -error of the SMC approximation

For each block n , denote by $\{v_{t,n}\}_{t \leq \tau_{n+1}}$ and $\{q_{t,n}\}_{t \leq \tau_{n+1}}$ respectively the adjustment multipliers and the instrumental kernels in the SMC propagation step (see (13)). For all $y \in \mathbb{Y}$, define

$$\omega_+(y) = \sup_{\theta \in \Theta} \sup_{\substack{(x, x') \in \mathbb{X} \times \mathbb{X} \\ t \geq 0, n \geq 0}} \frac{m_\theta(x, x') g_\theta(x', y)}{v_{t,n}(x) q_{t,n}(x, x')}.$$

Consider the following assumptions.

$$\mathbf{A6} \quad |v|_\infty \stackrel{\text{def}}{=} \sup_{t,n} |v_{t,n}|_\infty < \infty.$$

$$\mathbf{A7}\text{-(}p\text{)} \left\| \frac{\omega_+(\mathbf{Y}_0)}{b_-(\mathbf{Y}_0)} \right\|_p < +\infty.$$

In the case of the *Bootstrap filter*, A6 holds (since $v_{t,n} = 1$) and $\omega_+(y) = \sup_{\theta \in \Theta} \sup_{x \in \mathbb{X}} g_\theta(x, y)$.

Proposition 4.1. *Let $S : \mathbb{X}^2 \times \mathbb{Y} \longrightarrow \mathbb{R}^d$ be a measurable function s.t. A3-($\bar{p}$) holds for some $\bar{p} > 2$. Assume A2, A4, A6. Define $\Delta p \stackrel{\text{def}}{=} 2\bar{p}p/(\bar{p} - p)$ and assume A7-(Δp) holds for some $p \in (2, \bar{p})$. Then, there exists a constant C s.t. for all $n \geq 0$,*

$$\left\| \tilde{S}_n(\theta_n, \mathbf{Y}) - \bar{S}_{\tau_{n+1}}^{T_n}(\theta_n, \mathbf{Y}) \right\|_p \leq C \left(\frac{1}{N_{n+1}} + \frac{1}{\tau_{n+1}^{1/2} N_{n+1}^{1/2}} \right),$$

where $\tilde{S}_n(\theta_n, \mathbf{Y})$ is computed with the algorithm described in Section 2.4.

4.3 Asymptotic behavior of the Particle Block Online EM algorithms

Following [23], we address the convergence of the P-BOEM algorithm as the convergence of a perturbed version of the *limiting EM* recursion. The following result, which is proved in [23, Theorem 4.1.], shows that when τ is large, the BOEM statistic $\bar{S}_\tau^T(\theta, \mathbf{Y})$ is an approximation of a deterministic quantity $\bar{S}(\theta)$; the *limiting EM* algorithm is the iterative procedure defined by $\theta_{n+1} = R(\theta_n)$ where

$$R(\theta) \stackrel{\text{def}}{=} \bar{\theta}(\bar{S}(\theta)), \forall \theta \in \Theta; \quad (20)$$

the mapping $\bar{\theta}$ is given by A1.

Theorem 4.2. *Let $S : \mathbb{X}^2 \times \mathbb{Y} \longrightarrow \mathbb{R}^d$ be a measurable function s.t. A3-(1) holds. Assume A2 and A4(a). For any $\theta \in \Theta$, there exists a $\mathbb{P}$ -integrable r.v. $\mathbb{E}_\theta[S(X_{-1}, X_0, \mathbf{Y}_0)|\mathbf{Y}]$ s.t. for any $T > 0$,*

$$\bar{S}_\tau^T(\theta, \mathbf{Y}) \xrightarrow{\tau \rightarrow +\infty} \bar{S}(\theta) \stackrel{\text{def}}{=} \mathbb{E}[\mathbb{E}_\theta[S(X_{-1}, X_0, \mathbf{Y}_0)|\mathbf{Y}]], \quad \mathbb{P} - \text{a.s.} \quad (21)$$

Moreover, $\theta \mapsto \bar{S}(\theta)$ is continuous on Θ .

The asymptotic behavior of the limiting EM algorithm is addressed in [23, Section 4.2]: the main ingredient is that the map R admits a positive and continuous Lyapunov function W w.r.t. the set

$$\mathcal{L} \stackrel{\text{def}}{=} \{\theta \in \Theta; R(\theta) = \theta\}, \quad (22)$$

i.e. (i) $W \circ R(\theta) \geq W(\theta)$ for any $\theta \in \Theta$ and, (ii) for any compact subset $\mathcal{K}$ of $\Theta \setminus \mathcal{L}$, $\inf_{\theta \in \mathcal{K}} W \circ R(\theta) - W(\theta) > 0$. This Lyapunov function is equal to $\exp(\ell(\theta))$, where the contrast function $\ell(\theta)$ is the (deterministic) limit of the normalized log-likelihood of the observations when $T \rightarrow +\infty$ (see [24, Theorem 4.9]).

Theorem 4.3 establishes the convergence of the P-BOEM algorithm to the set $\mathcal{L}$ defined by (22). The proof of Theorem 4.3 is an application of [23, Theorem 4.4]. An additional assumption on the number of particles per block is required to check [23, A6] (note indeed that A8 below and Proposition 4.1 imply the condition in [23] about the L_p -control of the error).

A8 There exist $c > 0$ and $d \geq (a + 1)/2a$ (where a is given by A5) such that, for all $n \geq 1$, $N_n = \lfloor c\tau_n^d \rfloor$.

Theorem 4.3. Assume A1-2, A3-($\bar{p}$), A4-6 and A8 for some $\bar{p} > 2$. Define $\Delta p \stackrel{\text{def}}{=} 2\bar{p}p/(\bar{p}-p)$ and assume A7-(Δp) holds for some $p \in (2, \bar{p})$. Assume in addition that $W(\mathcal{L})$ has an empty interior. Then, there exists $w_\star$ s.t. $\{W(\theta_n)\}_{n \geq 0}$ converges almost surely to $w_\star$ and $\{\theta_n\}_{n \geq 0}$ converges to $\{\theta \in \mathcal{L}; W(\theta) = w_\star\}$.

The assumption on $W(\mathcal{L})$ made in Theorem 4.3 is in common use to prove the convergence of EM based procedures or stochastic approximation algorithms. It is used in [36] to find the limit points of the classical EM algorithm. See also [13] and [20] for the stability of the Monte Carlo EM algorithm and of a stochastic approximation of the EM algorithm. If W is sufficiently regular, Sard's theorem states that $W(\mathcal{L})$ has Lebesgue measure 0 and hence has an empty interior.

Under the assumptions of Theorem 4.3, it can be proved that, along any converging P-BOEM sequence $\{\theta_n\}_{n \geq 0}$ to $\theta_\star$ in $\mathcal{L}$, the averaged P-BOEM statistics $\{\tilde{\Sigma}_n\}_n$ defined by (7) (see also (9)) converge to $\bar{S}(\theta_\star)$, see Proposition 5.2. Since $\bar{\theta}$ is continuous, the averaged P-BOEM sequence $\{\tilde{\theta}_n\}_{n \geq 0}$ converges to $\bar{\theta}(\bar{S}(\theta_\star)) = R(\theta_\star)$. Since $\theta_\star \in \mathcal{L}$, $R(\theta_\star) = \theta_\star$, showing that the averaged P-BOEM algorithm has the same limit points as the P-BOEM algorithm.

4.4 Rate of convergence of the Particle Block Online EM algorithms

In this section, we consider a converging P-BOEM sequence $\{\theta_n\}_{n \geq 0}$ with limiting point $\theta_\star \in \mathcal{L}$. It can be shown, as in [24, Proposition 3.1], that the convergence of the sequence $\{\theta_n\}_{n \geq 0}$ is equivalent to the convergence of the sufficient statistics $\{\tilde{S}_n(\theta_n, \mathbf{Y})\}_{n \geq 0}$: along any P-BOEM sequence converging to $\theta_\star$, this sequence of sufficient statistics converges to $s_\star = \bar{S}(\theta_\star)$. Let $G : \mathcal{S} \rightarrow \mathcal{S}$ be the limiting EM map defined on the space of sufficient statistics by

$$G(s) \stackrel{\text{def}}{=} \bar{S}(\bar{\theta}(s)) , \quad \forall s \in \mathcal{S} . \quad (23)$$

To that goal consider the following assumption.

- A9** (a) $\bar{S}$ and $\bar{\theta}$ are twice continuously differentiable on Θ and $\mathcal{S}$.
(b) $\text{sp}(\nabla_s G(s_\star)) \in (0, 1)$ where sp denotes the spectral radius.

We will use the following notation: for any sequence of random variables $\{Z_n\}_{n \geq 0}$, write $Z_n = O_{L_p}(1)$ if $\limsup_n \|Z_n\|_p < \infty$; and $Z_n = O_{\text{a.s.}}(1)$ if $\sup_n |Z_n| < +\infty$ $\mathbb{P}$ -a.s.

Theorem 4.4. Assume A1-2, A3-($\bar{p}$), A4-6 and A8-9 for some $\bar{p} > 2$. Define $\Delta p \stackrel{\text{def}}{=} 2\bar{p}p/(\bar{p}-p)$ and assume A7-(Δp) holds for some $p \in (2, \bar{p})$. Then,

$$[\theta_n - \theta_\star] \mathbf{1}_{\lim_k \theta_k = \theta_\star} = O_{L_{p/2}} \left(\frac{1}{\tau_n^{1/2}} \right) O_{\text{a.s.}}(1) . \quad (24)$$

On the other hand, for the averaged sequence,

$$\left[\tilde{\theta}_n - \theta_\star \right] \mathbf{1}_{\lim_k \theta_k = \theta_\star} = O_{L_{p/2}} \left(\frac{1}{T_n^{1/2}} \right) O_{\text{a.s.}}(1) . \quad (25)$$

The proof of Theorem 4.4 is obtained by checking the assumptions of [23, Theorem 5.1 and Theorem 5.2].

Eq. (24) shows that the error $\theta_n - \theta_\star$ has a $L_{p/2}$ -norm decreasing as $\tau_n^{-1/2}$. This result is obtained by assuming $N_n \sim \tau_n^d$, with $d \geq (a+1)/2a$, which implies that the SMC error and the BOEM error are balanced. Unfortunately, such a rate is obtained after a total number of observations T_n ; therefore, as discussed in [23], it is quite sub-optimal. Eq (25) shows that the rate of convergence equal to the square root of the total number of observations up to block n , can be reached by using the averaged P-BOEM algorithm: the $L_{p/2}$ -norm of the error $\tilde{\theta}_n - \theta_\star$ has a rate of convergence proportional to $T_n^{-1/2}$. Here again, note that since N_n is chosen as in A8 the SMC error and the BOEM error are balanced.

5 Proofs

For a function h , define $\text{osc}(h) \stackrel{\text{def}}{=} \sup_{z, z'} |h(z) - h(z')|$.

5.1 Proof of Proposition 4.1

For any $t \in \{0, \dots, \tau_{n+1}\}$, define the σ -algebra $\mathcal{F}_{n,t}^{N_{n+1}}$ by

$$\mathcal{F}_{n,t}^{N_{n+1}} \stackrel{\text{def}}{=} \sigma \left\{ \theta_n, \mathbf{Y}_{T_n+1:T_n+t+1}, (\xi_s^\ell, \omega_s^\ell); \ell \in \{1, \dots, N_{n+1}\}; 0 \leq s \leq t \right\} . \quad (26)$$

We use $S_s(x, x')$ as a shorthand notation for $S(x, x', \mathbf{Y}_s)$. Under A2 and A6, Propositions B.5., B.8. and B.9. in Appendix B of [22] can be applied so that

$$\mathbb{E} \left[\left| \tilde{S}_n(\theta_n, \mathbf{Y}) - \bar{S}_{\tau_{n+1}}^{T_n}(\theta_n, \mathbf{Y}) \right|^p \right] \leq C (I_{1,n} + I_{2,n}) , \quad (27)$$

where

$$\begin{aligned} I_{1,n} &\stackrel{\text{def}}{=} \frac{1}{\tau_{n+1}^{\frac{p}{2}+1} N_{n+1}^{\frac{p}{2}}} \times \sum_{t=0}^{\tau_{n+1}} \mathbb{E} \left[\left| \frac{\omega_+(\mathbf{Y}_{t+T_n})}{b_-(\mathbf{Y}_{t+T_n})} \sum_{s=1}^{\tau_{n+1}} \rho^{|t-s|} \text{osc}\{S_{s+T_n}\} \right|^p \right] , \\ I_{2,n} &\stackrel{\text{def}}{=} \frac{1}{\tau_{n+1} N_{n+1}^p} \\ &\quad \times \sum_{t=0}^{\tau_{n+1}} \mathbb{E} \left[\left| \frac{\omega_+(\mathbf{Y}_{t+T_n})}{b_-(\mathbf{Y}_{t+T_n})} \right|^{2p} \mathbb{E} \left[\left| \sum_{s=1}^{\tau_{n+1}} \rho^{|t-s|} \text{osc}\{S_{s+T_n}\} \right|^{\bar{p}} \middle| \mathcal{F}_{n,t-1}^{N_{n+1}} \right]^{p/\bar{p}} \right] . \end{aligned}$$

By the Hölder inequality applied with $\alpha \stackrel{\text{def}}{=} \bar{p}/p \geq 1$ and $\beta^{-1} \stackrel{\text{def}}{=} 1 - \alpha^{-1}$,

$$I_{1,n} \leq \frac{1}{\tau_{n+1}^{\frac{p}{2}+1} N_{n+1}^{\frac{p}{2}}} \sum_{t=1}^{\tau_{n+1}} \left\| \sum_{s=1}^{\tau_{n+1}} \rho^{|t-s|} \text{osc}\{S_{s+T_n}\} \right\|_{\bar{p}}^p \times \left\| \frac{\omega_+(\mathbf{Y}_{t+T_n})}{b_-(\mathbf{Y}_{t+T_n})} \right\|_{2\bar{p}p/(\bar{p}-p)}^{2p} .$$

By A2, A3-($\bar{p}$), A4(a) and A7-(Δp), we have

$$I_{1,n} \leq \frac{C}{\tau_{n+1}^{\frac{p}{2}} N_{n+1}^{\frac{p}{2}}} .$$

Using similar arguments for $I_{2,n}$ yields $I_{2,n} \leq C N_{n+1}^{-p}$, which concludes the proof.

5.2 L_p -controls

Proposition 5.1. *Let $S : \mathbb{X}^2 \times \mathbb{Y} \rightarrow \mathbb{R}^d$ be a measurable function s.t. A3-($\bar{p}$) holds for some $\bar{p} > 2$. Assume A2, A4, A6 and A7-(Δp) for some $p \in (2, \bar{p})$, where $\Delta p \stackrel{\text{def}}{=} 2\bar{p}p/(\bar{p} - p)$. There exists a constant C s.t. for any $n \geq 1$,*

$$\left\| \tilde{S}_n(\theta_n, \mathbf{Y}) - \bar{S}(\theta_n) \right\|_p \leq C \left(\frac{1}{\sqrt{\tau_{n+1}}} + \frac{1}{N_{n+1}} \right) .$$

Proof. Under A2, A3-($\bar{p}$) and A4, by [23, Theorem 4.1], there exists a constant C s.t.

$$\left\| \bar{S}_{\tau_{n+1}}^{T_n}(\theta_n, \mathbf{Y}) - \bar{S}(\theta_n) \right\|_p \leq \frac{C}{\sqrt{\tau_{n+1}}} .$$

Moreover, under A6 and A7-(Δp), by Proposition 4.1, we have

$$\left\| \tilde{S}_n(\theta_n, \mathbf{Y}) - \bar{S}_{\tau_{n+1}}^{T_n}(\theta_n, \mathbf{Y}) \right\|_p \leq C \left(\frac{1}{N_{n+1}} + \frac{1}{\tau_{n+1}^{1/2} N_{n+1}^{1/2}} \right) ,$$

which concludes the proof. $\square$

Proposition 5.2. *Let $S : \mathbb{X}^2 \times \mathbb{Y} \rightarrow \mathbb{R}^d$ be a measurable function s.t. A3-($\bar{p}$) holds for some $\bar{p} > 2$. Assume A2, A4-5, A6-8 and A7-(Δp) for some $p \in (2, \bar{p})$, where $\Delta p \stackrel{\text{def}}{=} 2\bar{p}p/(\bar{p} - p)$. Let $\{\theta_n\}_n$ be the P-BOEM sequence. For any $\theta_\star \in \Theta$, on the set $\{\lim_n \theta_n = \theta_\star\}$,*

$$\tilde{\Sigma}_n \rightarrow \bar{S}(\theta_\star) , \quad \mathbb{P} - \text{a.s.} ,$$

where $\bar{S}$ is defined in (21) and $\tilde{\Sigma}_n$ in (7).

Proof. By (7), $\tilde{\Sigma}_n$ can be written as

$$\tilde{\Sigma}_n = \frac{1}{T_n} \sum_{j=1}^n \tau_j \left[\tilde{S}_{j-1}(\theta_{j-1}, \mathbf{Y}) - \bar{S}(\theta_{j-1}) \right] + \frac{1}{T_n} \sum_{j=1}^n \tau_j \bar{S}(\theta_{j-1}) . \quad (28)$$

By Theorem 4.2, $\bar{S}$ is continuous so, by the Cesaro Lemma, the second term in the right-hand side of (28) converges to $\bar{S}(\theta_\star)$ $\mathbb{P}$ -a.s., on the set $\{\lim_n \theta_n = \theta_\star\}$. By Proposition 5.1, there exists a constant C such that for any n ,

$$\left\| \tilde{S}_n(\theta_n, \mathbf{Y}) - \bar{S}(\theta_n) \right\|_p \leq C \left(\frac{1}{\sqrt{\tau_{n+1}}} + \frac{1}{N_{n+1}} \right) .$$

Hence, by A5, A8 and the Borel-Cantelli Lemma,

$$\left| \tilde{S}_n(\theta_n, \mathbf{Y}) - \bar{S}(\theta_n) \right| \longrightarrow 0, \quad \mathbb{P} - \text{a.s.}$$

The proof is concluded by applying the Cesaro Lemma. $\square$

A Detailed SMC algorithm

In this section, we give a detailed description of the SMC algorithm used to compute sequentially the quantities $\tilde{S}_n(\theta_n, \mathbf{Y})$, $n \geq 0$. This is the algorithm proposed by [4] and [10].

At each time step, the weighted samples are produced using sequential importance sampling and sampling importance resampling steps. In Algorithm 2, the instrumental proposition kernel used to select and propagate the particles is π_t (see (13) and [16, 17, 27] for further details on this SMC step).

It is readily seen from the description below that the observations $\mathbf{Y}_t$ are processed sequentially.

Algorithm 2 Forward SMC step

Require: $\theta_n, \tau_{n+1}, N, \mathbf{Y}_{T_n+1:T_n+\tau_{n+1}}$.

Ensure: $\tilde{S}_n(\theta_n, \mathbf{Y})$.

Sample $\{\xi_0^\ell\}_{\ell=1}^N$ i.i.d. with distribution χ .

Set $\omega_0^\ell = 1/N$ for all $\ell \in \{1, \dots, N\}$.

Set $R_{0,\theta_n}^\ell = 0$ for all $\ell \in \{1, \dots, N\}$.

for $t = 1$ to τ_{n+1} **do**

for $\ell = 1$ to N **do**

 Conditionally to $(\theta_n, Y_{T_n+1:T_n+t}, \{J_{t-1}^\ell, \xi_{t-1}^\ell\}_{\ell=1}^N)$, sample independently $(J_t^\ell, \xi_t^\ell) \sim \pi_t(i, dx)$, where $\pi_t(i, dx) \propto \omega_{t-1}^i v_t(\xi_{t-1}^i) q_t(\xi_{t-1}^i, x) \lambda(dx)$.

 Set

$$\omega_t^\ell = \frac{m_{\theta_n}(\xi_{t-1}^{J_t^\ell}, \xi_t^\ell) g_{\theta_n}(\xi_t^\ell, \mathbf{Y}_{T_n+t})}{v_t(\xi_{t-1}^{J_t^\ell}) q_t(\xi_{t-1}^{J_t^\ell}, \xi_t^\ell)}.$$

 Set

$$R_{t,\theta_n}^\ell = \frac{1}{t} \sum_{j=1}^N \omega_{t-1}^j m_{\theta_n}(\xi_{t-1}^j, \xi_t^\ell) \frac{S(\xi_{t-1}^j, \xi_t^\ell, \mathbf{Y}_{T_n+t}) + (t-1) R_{t-1,\theta_n}^j}{\sum_{k=1}^N \omega_{t-1}^k m_{\theta_n}(\xi_{t-1}^k, \xi_t^\ell)}.$$

end for

end for

Set

$$\tilde{S}_n(\theta_n, \mathbf{Y}) = \sum_{\ell=1}^N \omega_{\tau_{n+1}}^\ell R_{\tau_{n+1},\theta_n}^\ell.$$

B L_p -controls of SMC approximations

In this section, we give further details on the L_p control on each block (see (27)):

$$\mathbb{E} \left[\left| \tilde{S}_{\tau_{n+1}}^{N, T_n}(\theta_n, \mathbf{Y}) - \bar{S}_{\tau_{n+1}}^{T_n}(\theta_n, \mathbf{Y}) \right|^p \right],$$

$\bar{S}_\tau^T$ is defined by (6) (we recall that, χ being fixed, it is dropped from the notations) and $\tilde{S}_\tau^{N, T}$ is the SMC approximation of $\bar{S}_\tau^T$ based on N particles computed as described in Section 2.4.

The following results are technical lemmas taken from [16] (stated here for a better clarity) or extensions of the L_p controls derived in [19].

Hereafter, “time t ” corresponds to time t in the block n . Therefore, even if it is not explicit in the notations (in order to make them simpler), the following quantities depend upon the observations $\mathbf{Y}_{T_n+1:T_n+\tau_{n+1}}$.

Denote by ϕ_s^θ the filtering distribution at time s , and let

$$B_{\phi_t^\theta}^\theta(x, dx') \stackrel{\text{def}}{=} \frac{m_\theta(x', x)}{\int m_\theta(u, x) \phi_t^\theta(du)} \phi_t^\theta(dx')$$

be the backward kernel smoothing kernel at time $t+1$. For all $0 \leq s \leq \tau-1$ and for all bounded measurable function h on $\mathbb{X}^{\tau-s+1}$, define recursively $\phi_{s:\tau|\tau}^\theta[h]$ backward in time, according to

$$\phi_{s:\tau|\tau}^\theta[h] = \int \cdots \int B_{\phi_s^\theta}^\theta(x_{s+1}, dx_s) \phi_{s+1:\tau|\tau}^\theta(dx_{s+1:\tau}) h(x_{s:\tau}), \quad (29)$$

starting from $\phi_{\tau:\tau|\tau}^\theta = \phi_\tau^\theta$. By convention, $\phi_0^\theta = \chi$.

For $t \geq 1$, let $\{(\xi_t^\ell, \omega_t^\ell)\}_{\ell=1}^N$ be the weighted samples obtained as described in Section 2.4 (see also Algorithm 2 in Appendix A); it approximates the filtering distribution ϕ_t^θ . Denote by $\phi_t^{N, \theta}$ this approximation. For $0 \leq s \leq \tau-1$, an approximation of the backward kernel can be obtained

$$B_{\phi_s^{N, \theta}}(x, h) = \sum_{i=1}^N \frac{\omega_s^i m_\theta(\xi_s^i, x)}{\sum_{\ell=1}^N \omega_s^\ell m_\theta(\xi_s^\ell, x)} h(\xi_s^i);$$

and inserting this expression into (29) gives the following particle approximation of the fixed-interval smoothing distribution $\phi_{0:\tau|\tau}^\theta[h]$

$$\phi_{0:\tau|\tau}^{N, \theta}[h] = \sum_{i_0=1}^N \cdots \sum_{i_\tau=1}^N \left(\prod_{u=1}^{\tau} \frac{\omega_{u-1}^{i_u-1} m_\theta(\xi_{u-1}^{i_u-1}, \xi_u^{i_u})}{\sum_{\ell=1}^N \omega_{u-1}^\ell m_\theta(\xi_{u-1}^\ell, \xi_u^{i_u})} \right) \times \frac{\omega_\tau^{i_\tau}}{\Omega_\tau^n} h(\xi_0^{i_0}, \dots, \xi_\tau^{i_\tau}), \quad (30)$$

with $\Omega_\tau^N \stackrel{\text{def}}{=} \sum_{\ell=1}^N \omega_\tau^\ell$.

Lemma B.1. Let $\{(\xi_t^\ell, \omega_t^\ell), 1 \leq \ell \leq N, 0 \leq t \leq \tau_{n+1}\}$ be the weighted samples obtained by Algorithm 2 in Appendix A, with $\theta_n, \tau_{n+1}, N, \mathbf{Y}_{T_n+1:T_n+\tau_{n+1}}$. Then,

$$\begin{aligned} & \left[\tilde{S}_{\tau_{n+1}}^{N, T_n}(\theta_n, \mathbf{Y}) - \bar{S}_{\tau_{n+1}}^{T_n}(\theta_n, \mathbf{Y}) \right] \\ &= \frac{1}{\tau_{n+1}} \left(\phi_{0:\tau_{n+1}|\tau_{n+1}}^{N, \theta_n} [\mathbf{S}_{\tau_{n+1}}] - \phi_{0:\tau_{n+1}|\tau_{n+1}}^{\theta_n} [\mathbf{S}_{\tau_{n+1}}] \right), \end{aligned} \quad (31)$$

where

$$\mathbf{S}_\tau(x_{0:\tau}) \stackrel{\text{def}}{=} \sum_{s=1}^{\tau} S(x_{s-1}, x_s, \mathbf{Y}_{s:T}). \quad (32)$$

For all $t \in \{0, \dots, \tau\}$ and all bounded measurable function h on $\mathbb{X}^{\tau+1}$, define the kernel $\mathbf{L}_{t,\tau} : \mathbb{X}^{t+1} \times \mathcal{X}^{\otimes \tau+1} \rightarrow [0, 1]$ by

$$\mathbf{L}_{t,\tau}^\theta h(x_{0:t}) \stackrel{\text{def}}{=} \int \prod_{u=t+1}^{\tau} m_\theta(x_{u-1}, x_u) g_\theta(x_u, \mathbf{Y}_{u:T}) h(x_{0:\tau}) \lambda(dx_{t+1:\tau}); \quad (33)$$

by convention, $\mathbf{L}_{\tau,\tau}^\theta h = h$. Let $\mathcal{L}_{t,\tau}^{N,\theta}$ and $\mathcal{L}_{t,\tau}^\theta$ be two kernels on $\mathbb{X} \times \mathcal{X}^{\otimes(\tau+1)}$ defined for all $x_t \in \mathbb{X}$ by

$$\mathcal{L}_{t,\tau}^\theta h(x_t) \stackrel{\text{def}}{=} \int \mathbf{B}_{\phi_{t-1}^\theta}^\theta(x_t, dx_{t-1}) \cdots \mathbf{B}_{\phi_0^\theta}^\theta(x_1, dx_0) \mathbf{L}_{t,\tau}^\theta h(x_{0:t}) \quad (34)$$

$$\mathcal{L}_{t,\tau}^{N,\theta} h(x_t) \stackrel{\text{def}}{=} \int \mathbf{B}_{\phi_{t-1}^{N,\theta}}^\theta(x_t, dx_{t-1}) \cdots \mathbf{B}_{\phi_0^{N,\theta}}^\theta(x_1, dx_0) \mathbf{L}_{t,\tau}^\theta h(x_{0:t}). \quad (35)$$

Note that

$$\mathcal{L}_{t,\tau}^\theta \mathbf{1}(x_t) = \int m_\theta(x_t, x') g_\theta(x', \mathbf{Y}_{T+t+1}) \mathcal{L}_{t+1,\tau}^\theta \mathbf{1}(x') \lambda(dx'). \quad (36)$$

Lemma B.2, Proposition B.3, Lemma B.4 and B.5 can be found in [16].

Lemma B.2. Let $\{(\xi_t^\ell, \omega_t^\ell), 1 \leq \ell \leq N, 0 \leq t \leq \tau_{n+1}\}$ be the weighted samples obtained by Algorithm 2 in Appendix A, with $\theta_n, \tau_{n+1}, N, \mathbf{Y}_{T_n+1:T_n+\tau_{n+1}}$. Then,

$$\phi_{0:\tau_{n+1}|\tau_{n+1}}^{N, \theta_n} [h] - \phi_{0:\tau_{n+1}|\tau_{n+1}}^{\theta_n} [h] = \sum_{t=0}^{\tau_{n+1}} \frac{\sum_{\ell=1}^N \omega_t^\ell G_{t,\tau_{n+1}}^{N, \theta_n} h(\xi_t^\ell)}{\sum_{\ell=1}^N \omega_t^\ell \mathcal{L}_{t,\tau_{n+1}}^{\theta_n} \mathbf{1}(\xi_t^\ell)}, \quad (37)$$

with $G_{t,\tau}^{N,\theta}$ is a kernel on $\mathbb{X} \times \mathcal{X}^{\otimes(\tau+1)}$ defined, for all $x \in \mathbb{X}$ and all bounded and measurable function h on $\mathbb{X}^{\tau+1}$, by

$$G_{t,\tau}^{N,\theta} h(x) \stackrel{\text{def}}{=} \mathcal{L}_{t,\tau}^{N,\theta} h(x) - \frac{\phi_{t-1}^{N,\theta} [\mathcal{L}_{t-1,\tau}^{N,\theta} h]}{\phi_{t-1}^{N,\theta} [\mathcal{L}_{t-1,\tau}^{N,\theta} \mathbf{1}]} \mathcal{L}_{t,\tau}^{N,\theta} \mathbf{1}(x). \quad (38)$$

Proof. By definition of $L_{t,\tau}^\theta$,

$$\phi_{0:\tau|\tau}^\theta[h] = \frac{\phi_{0:t|t}^\theta[L_{t,\tau}^\theta h]}{\phi_{0:t|t}^\theta[L_{t,\tau}^\theta \mathbf{1}]}.$$

We write

$$\phi_{0:\tau|\tau}^{N,\theta}[h] - \phi_{0:\tau|\tau}^\theta[h] = \sum_{t=0}^{\tau} \left\{ \frac{\phi_{0:t|t}^{N,\theta}[L_{t,\tau}^\theta h]}{\phi_{0:t|t}^{N,\theta}[L_{t,\tau}^\theta \mathbf{1}]} - \frac{\phi_{0:t-1|t-1}^{N,\theta}[L_{t-1,\tau}^\theta h]}{\phi_{0:t-1|t-1}^{N,\theta}[L_{t-1,\tau}^\theta \mathbf{1}]} \right\},$$

where we used the convention

$$\frac{\phi_{0:-1|t-1}^{N,\theta}[L_{t-1,\tau}^\theta h]}{\phi_{0:-1|t-1}^{N,\theta}[L_{t-1,\tau}^\theta \mathbf{1}]} = \frac{\chi[L_{0,\tau}^\theta h]}{\chi[L_{0,\tau}^\theta \mathbf{1}]} = \phi_{0:\tau|\tau}^\theta[h].$$

We have for all $0 \leq t \leq \tau$,

$$\phi_{0:t|t}^{N,\theta}[L_{t,\tau}^\theta h] = \int \phi_t^{N,\theta}(dx_t) \prod_{j=0}^{t-1} B_{\phi_j^{N,\theta}}(x_{j+1}, dx_j) L_{t,\tau}^\theta h(x_{0:t}) = \phi_t^{N,\theta}[L_{t,\tau}^\theta h].$$

Therefore, for all $1 \leq t \leq \tau$,

$$\frac{\phi_{0:t|t}^{N,\theta}[L_{t,\tau}^\theta h]}{\phi_{0:t|t}^{N,\theta}[L_{t,\tau}^\theta \mathbf{1}]} - \frac{\phi_{0:t-1|t-1}^{N,\theta}[L_{t-1,\tau}^\theta h]}{\phi_{0:t-1|t-1}^{N,\theta}[L_{t-1,\tau}^\theta \mathbf{1}]} = \frac{\phi_t^{N,\theta}[L_{t,\tau}^\theta h]}{\phi_t^{N,\theta}[L_{t,\tau}^\theta \mathbf{1}]} - \frac{\phi_{t-1}^{N,\theta}[L_{t-1,\tau}^\theta h]}{\phi_{t-1}^{N,\theta}[L_{t-1,\tau}^\theta \mathbf{1}]} = \frac{\phi_t^{N,\theta}[G_{t,\tau}^\theta h]}{\phi_t^{N,\theta}[L_{t,\tau}^\theta \mathbf{1}]}.$$

□

Proposition B.3. Let $\{(\xi_t^\ell, \omega_t^\ell), 1 \leq \ell \leq N, 0 \leq t \leq \tau_{n+1}\}$ be the weighted samples obtained by Algorithm 2 in Appendix A, with input variables $\theta_n, \tau_{n+1}, N, \mathbf{Y}_{T_n+1:T_n+\tau_{n+1}}$. Then,

$$\begin{aligned} \left[\tilde{S}_{\tau_{n+1}}^{N,T_n}(\theta_n, \mathbf{Y}) - \bar{S}_{\tau_{n+1}}^{T_n}(\theta_n, \mathbf{Y}) \right] &= \frac{1}{\tau_{n+1}} \sum_{t=0}^{\tau_{n+1}} D_{t,\tau_{n+1}}^{N,\theta_n}(S_{\tau_{n+1}}) \\ &\quad + \frac{1}{\tau_{n+1}} \sum_{t=0}^{\tau_{n+1}} C_{t,\tau_{n+1}}^{N,\theta_n}(S_{\tau_{n+1}}), \end{aligned}$$

where S_τ is given by (32) and

$$D_{t,\tau_{n+1}}^{N,\theta_n}(h) \stackrel{\text{def}}{=} \frac{\phi_{t-1}^{N,\theta_n}[v_t]}{\phi_{t-1}^{N,\theta_n} \left[\frac{\mathcal{L}_{t-1,\tau_{n+1}}^{\theta_n} \mathbf{1}}{|\mathcal{L}_{t,\tau_{n+1}}^{\theta_n} \mathbf{1}|_\infty} \right]} N^{-1} \sum_{\ell=1}^N \omega_t^\ell \frac{G_{t,\tau_{n+1}}^{N,\theta_n} h(\xi_t^\ell)}{|\mathcal{L}_{t,\tau_{n+1}}^{\theta_n} \mathbf{1}|_\infty}; \quad (39)$$

$$\begin{aligned} C_{t,\tau_{n+1}}^N(h) &\stackrel{\text{def}}{=} \left[\frac{1}{N^{-1} \sum_{i=1}^N \omega_t^i \frac{\mathcal{L}_{t,\tau_{n+1}}^{\theta_n} \mathbf{1}(\xi_t^i)}{|\mathcal{L}_{t,\tau_{n+1}}^{\theta_n} \mathbf{1}|_\infty}} - \frac{\phi_{t-1}^{N,\theta_n}[v_t]}{\phi_{t-1}^{N,\theta_n} \left[\frac{\mathcal{L}_{t-1,\tau_{n+1}}^{\theta_n} \mathbf{1}}{|\mathcal{L}_{t,\tau_{n+1}}^{\theta_n} \mathbf{1}|_\infty} \right]} \right] \\ &\quad \times N^{-1} \sum_{\ell=1}^N \omega_t^\ell \frac{G_{t,\tau_{n+1}}^{N,\theta_n} h(\xi_t^\ell)}{|\mathcal{L}_{t,\tau_{n+1}}^{\theta_n} \mathbf{1}|_\infty}. \end{aligned} \quad (40)$$

Proof. (37) can be rewritten as follows:

$$\phi_{0:\tau_{n+1}|\tau_{n+1}}^{N,\theta_n}[h] - \phi_{0:\tau_{n+1}|\tau_{n+1}}^{\theta_n}[h] = \sum_{t=0}^{\tau_{n+1}} D_{t,\tau_{n+1}}^{N,\theta_n}(h) + \sum_{t=0}^{\tau_{n+1}} C_{t,\tau_{n+1}}^{N,\theta_n}(h). \quad (41)$$

The proof is concluded by Lemma B.1. $\square$

For any $t \in \{0, \dots, \tau_{n+1}\}$, we recall the definition of $\mathcal{F}_{n,t}^N$ given by (26)

$$\mathcal{F}_{n,t}^N = \sigma \left\{ \theta_n, \mathbf{Y}_{T_n+1:T_n+t+1}, (\xi_s^\ell, \omega_s^\ell); \ell \in \{1, \dots, N\}; 0 \leq s \leq t \right\},$$

where $\{(\xi_t^\ell, \omega_t^\ell)\}_{\ell=1}^N$ are the weighted samples obtained by Algorithm 2 in Appendix A, with input variables $\theta_n, \tau_{n+1}, N, \mathbf{Y}_{T_n+1:T_n+\tau_{n+1}}$.

Lemma B.4. *Let $\{(\xi_t^\ell, \omega_t^\ell), 1 \leq \ell \leq N, 0 \leq t \leq \tau_{n+1}\}$ be the weighted samples obtained by Algorithm 2 in Appendix A, with $\theta_n, \tau_{n+1}, N, \mathbf{Y}_{T_n+1:T_n+\tau_{n+1}}$. Then, for any $1 \leq t \leq \tau_{n+1}$ and any $1 \leq \ell \leq N$,*

$$\mathbb{E} [\omega_t^\ell h(\xi_t^\ell) | \mathcal{F}_{n,t-1}^N] = \frac{\phi_{t-1}^{N,\theta_n} [\int m_{\theta_n}(\cdot, x) g_{\theta_n}(x, \mathbf{Y}_{T_n+t}) h(x) \lambda(dx)]}{\phi_{t-1}^{N,\theta_n}[v_t]}. \quad (42)$$

Proof. By definition of the weighted particles,

$$\begin{aligned} & \mathbb{E} [\omega_t^\ell h(\xi_t^\ell) | \mathcal{F}_{n,t-1}^N] \\ &= \mathbb{E} \left[\frac{m_{\theta_n}(\xi_{t-1}^{I_t^1}, \xi_t^1) g_{\theta_n}(\xi_t^1, \mathbf{Y}_{t+T_n})}{v_t(\xi_{t-1}^{I_t^1}) q_t(\xi_{t-1}^{I_t^1}, \xi_t^1)} h(\xi_t^1) \middle| \mathcal{F}_{n,t-1}^N \right] \\ &= \left(\sum_{i=1}^N \omega_{t-1}^i v_t(\xi_{t-1}^i) \right)^{-1} \sum_{i=1}^N \int \omega_{t-1}^i v_t(\xi_{t-1}^i) q_t(\xi_{t-1}^i, x) \\ & \quad \times \frac{m_{\theta_n}(\xi_{t-1}^i, x) g_{\theta_n}(x, \mathbf{Y}_{t+T_n})}{v_t(\xi_{t-1}^i) q_t(\xi_{t-1}^i, x)} h(x) \lambda(dx) \\ &= \left(\sum_{i=1}^N \omega_{t-1}^i v_t(\xi_{t-1}^i) \right)^{-1} \sum_{i=1}^N \int \omega_{t-1}^i m_{\theta_n}(\xi_{t-1}^i, x) g_{\theta_n}(x, \mathbf{Y}_{t+T_n}) h(x) \lambda(dx). \end{aligned}$$

$\square$

Lemma B.5. *Assume A2 and A6. Let $\{(\xi_t^\ell, \omega_t^\ell), 1 \leq \ell \leq N, 0 \leq t \leq \tau_{n+1}\}$ be the weighted samples obtained by Algorithm 2 in Appendix A, with input variables $\theta_n, \tau_{n+1}, N, \mathbf{Y}_{T_n+1:T_n+\tau_{n+1}}$.*

(i) *For any $t \in \{0, \dots, \tau_{n+1}\}$ and any measurable function h on $\mathbb{X}^{\tau_{n+1}+1}$, the random variables $\left\{ \omega_t^\ell G_{t,\tau_{n+1}}^{N,\theta_n} h(\xi_t^\ell) \mid \mathcal{L}_{t,\tau_{n+1}}^{\theta_n} \mathbf{1}_{\infty}^{-1} \right\}_{\ell=1}^N$ are:*

(a) *conditionally independent and identically distributed given $\mathcal{F}_{n,t-1}^N$,*

(b) centered conditionally to $\mathcal{F}_{n,t-1}^N$.

(ii) For any $t \in \{0, \dots, \tau_{n+1}\}$:

$$\left| \frac{G_{t,\tau_{n+1}}^{N,\theta_n} \mathbf{S}_{\tau_{n+1}}(\xi_t^\ell)}{|\mathcal{L}_{t,\tau_{n+1}}^{\theta_n} \mathbf{1}|_\infty} \right| \leq \sum_{s=1}^{\tau_{n+1}} \rho^{|t-s|} \text{osc}\{S(\cdot, \cdot, \mathbf{Y}_{s+T_n})\}, \quad (43)$$

where $\mathbf{S}_\tau$ is defined by (32).

(iii) For all $x \in \mathbb{X}$ and any $t \in \{0, \dots, \tau_{n+1}\}$,

$$\frac{\mathcal{L}_{t,\tau_{n+1}}^{\theta_n} \mathbf{1}(x)}{|\mathcal{L}_{t,\tau_{n+1}}^{\theta_n} \mathbf{1}|_\infty} \geq \frac{\sigma_-}{\sigma_+}, \quad \frac{\mathcal{L}_{t-1,\tau_{n+1}}^{\theta_n} \mathbf{1}(x)}{|\mathcal{L}_{t,\tau_{n+1}}^{\theta_n} \mathbf{1}|_\infty} \geq \frac{\sigma_-^2}{\sigma_+} b_-(\mathbf{Y}_{t+T_n}).$$

Proof. The proof of (i) is given by [16, Lemma 3].

Proof of (ii). Let $\Pi_{s-1:s,\tau}$ be the operator which associates to any bounded and measurable function h on $\mathbb{X} \times \mathbb{X}$ the function $\Pi_{s-1:s,\tau} h$ given, for any $(x_0, \dots, x_\tau) \in \mathbb{X}^{\tau+1}$, by

$$\Pi_{s-1:s,\tau} h(x_{0:\tau}) \stackrel{\text{def}}{=} h(x_{s-1:s}).$$

Using this notation, we may write $\mathbf{S}_\tau = \sum_{s=1}^\tau \Pi_{s-1:s,\tau} S(\cdot, \cdot, \mathbf{Y}_{s+T})$ and $G_{t,\tau}^{N,\theta} \mathbf{S}_\tau = \sum_{s=1}^\tau G_{t,\tau}^{N,\theta} \Pi_{s-1:s,\tau} S(\cdot, \cdot, \mathbf{Y}_{s+T})$. Following the same lines as in [16, Lemma 10],

$$\begin{aligned} |G_{t,\tau}^{N,\theta} \Pi_{s-1:s,\tau} S(\cdot, \cdot, \mathbf{Y}_{s+T})|_\infty &\leq \rho^{s-1-t} \text{osc}(S(\cdot, \cdot, \mathbf{Y}_{s+T})) |\mathcal{L}_{t,\tau}^\theta \mathbf{1}|_\infty \quad \text{if } t \leq s-1, \\ |G_{t,\tau}^{N,\theta} \Pi_{s-1:s,\tau} S(\cdot, \cdot, \mathbf{Y}_{s+T})|_\infty &\leq \rho^{t-s} \text{osc}(S(\cdot, \cdot, \mathbf{Y}_{s+T})) |\mathcal{L}_{t,\tau}^\theta \mathbf{1}|_\infty \quad \text{if } t \geq s. \end{aligned}$$

Consequently,

$$\begin{aligned} \left| G_{t,\tau}^{N,\theta} \mathbf{S}_\tau \right|_\infty &\leq \sum_{s=1}^\tau |G_{t,\tau}^{N,\theta} \Pi_{s-1:s,\tau} S(\cdot, \cdot, \mathbf{Y}_{s+T})|_\infty \\ &\leq \left(\sum_{s=1}^\tau \rho^{|t-s|} \text{osc}\{S(\cdot, \cdot, \mathbf{Y}_{s+T})\} \right) |\mathcal{L}_{t,\tau}^\theta \mathbf{1}|_\infty, \end{aligned}$$

which shows (ii).

Proof of (iii). By the definition (34), for all $x \in \mathbb{X}$ and all $t \in \{1, \dots, \tau\}$,

$$\begin{aligned} \mathcal{L}_{t,\tau}^\theta \mathbf{1}(x) &= \int m_\theta(x, x_{t+1}) g_\theta(x_{t+1}, \mathbf{Y}_{t+T+1}) \\ &\quad \times \prod_{u=t+2}^\tau m_\theta(x_{u-1}, dx_u) g_\theta(x_u, \mathbf{Y}_{u+T}) \lambda(dx_{t+1:\tau}). \end{aligned}$$

Hence, by A2,

$$\begin{aligned} |\mathcal{L}_{t,\tau}^\theta \mathbf{1}|_\infty &\leq \sigma_+ \int g_\theta(x_{t+1}, \mathbf{Y}_{t+T+1}) \mathcal{L}_{t+1,\tau}^\theta \mathbf{1}(x_{t+1}) \lambda(dx_{t+1}) \\ \mathcal{L}_{t,\tau}^\theta \mathbf{1}(x) &\geq \sigma_- \int g_\theta(x_{t+1}, \mathbf{Y}_{t+T+1}) \mathcal{L}_{t+1,\tau}^\theta \mathbf{1}(x_{t+1}) \lambda(dx_{t+1}), \end{aligned}$$

which concludes the proof of the first statement. By (36), A2 and (16),

$$\frac{\mathcal{L}_{t-1,\tau}^\theta \mathbf{1}(x)}{|\mathcal{L}_{t,\tau}^\theta \mathbf{1}|_\infty} = \int m_\theta(x, x') g_\theta(x', \mathbf{Y}_{t+T}) \frac{\mathcal{L}_{t,\tau}^\theta \mathbf{1}(x')}{|\mathcal{L}_{t,\tau}^\theta \mathbf{1}|_\infty} \lambda(dx') \geq \frac{\sigma_-^2}{\sigma_+} b_-(\mathbf{Y}_{t+T}).$$

□

The proofs of Propositions B.6 and B.7 follow the same lines as [19, Propositions 1-2]. The upper bounds given here provide an explicit dependence on the observations.

Proposition B.6. *Assume A2 and A6. Let $\{(\xi_t^\ell, \omega_t^\ell), 1 \leq \ell \leq N, 0 \leq t \leq \tau_{n+1}\}$ be the weighted samples obtained by Algorithm 2 in Appendix A, with input variables $\theta_n, \tau_{n+1}, N, \mathbf{Y}_{T_{n+1}:T_n+\tau_{n+1}}$. For all $p > 1$, there exists a constant C such that*

$$\begin{aligned} &\mathbb{E} \left[\left| \sum_{t=0}^{\tau_{n+1}} D_{t,\tau_{n+1}}^{N,\theta_n} (S_{\tau_{n+1}}) \right|^p \right] \\ &\leq C \frac{\tau_{n+1}^{(\frac{p}{2}-1) \vee 0}}{N^{p-(\frac{p}{2} \vee 1)}} \sum_{t=0}^{\tau_{n+1}} \mathbb{E} \left[\left| \frac{\omega_+(\mathbf{Y}_{t+T_n})}{b_-(\mathbf{Y}_{t+T_n})} \sum_{s=1}^{\tau_{n+1}} \rho^{|t-s|} \text{osc}\{S(\cdot, \cdot, \mathbf{Y}_{s+T_n})\} \right|^p \right]. \quad (44) \end{aligned}$$

where $D_{t,\tau}^{N,\theta}$ is defined in (39).

Proof. By Lemma B.5(iii),

$$\frac{\phi_{t-1}^{N,\theta_n}[v_t]}{\phi_{t-1}^{N,\theta_n} \left[\frac{\mathcal{L}_{t-1,\tau_{n+1}}^{\theta_n} \mathbf{1}}{|\mathcal{L}_{t,\tau_{n+1}}^{\theta_n} \mathbf{1}|_\infty} \right]} \leq \frac{\sigma_+ |v|_\infty}{\sigma_-^2 b_-(\mathbf{Y}_{t+T_n})}.$$

By Lemma B.5(i) and since θ_n is $\mathcal{F}_{n,t}^N$ -measurable for all $t \in \{0, \dots, \tau_{n+1}\}$, $\left\{ D_{t,\tau_{n+1}}^{N,\theta_n} (S_{\tau_{n+1}}), \mathcal{F}_{n,t}^N \right\}_{0 \leq t \leq \tau_{n+1}}$ is a martingale difference. Since $p > 1$, Burkholder's inequality (see [21, Theorem 2.10, page 23]) states the existence of a constant C depending only on p such that:

$$\mathbb{E} \left[\left| \sum_{t=0}^{\tau_{n+1}} D_{t,\tau_{n+1}}^{N,\theta_n} (S_{\tau_{n+1}}) \right|^p \right] \leq C \mathbb{E} \left[\sum_{t=0}^{\tau_{n+1}} \left| D_{t,\tau_{n+1}}^{N,\theta_n} (S_{\tau_{n+1}}) \right|^2 \right]^{p/2}.$$

Hence,

$$\begin{aligned} \mathbb{E} \left[\left| \sum_{t=0}^{\tau_{n+1}} D_{t, \tau_{n+1}}^{N, \theta_n} (\mathbf{S}_{\tau_{n+1}}) \right|^p \right] &\leq C \left(\frac{\sigma_+ |v|_\infty}{\sigma_-^2} \right)^p \\ &\times \mathbb{E} \left[\left| \sum_{t=0}^{\tau_{n+1}} \left| N^{-1} \sum_{\ell=1}^N \frac{\omega_t^\ell}{b_-(\mathbf{Y}_{t+T_n})} \frac{G_{t, \tau_{n+1}}^{N, \theta_n} \mathbf{S}_{\tau_{n+1}}(\xi_t^\ell)}{|\mathcal{L}_{t, \tau_{n+1}}^{\theta_n} \mathbf{1}|_\infty} \right|^2 \right|^{p/2} \right], \end{aligned}$$

which implies, using the convexity inequality $(\sum_{k=1}^{\tau} a_k)^{p/2} \leq \tau^{(p/2-1)\vee 0} \sum_{k=1}^{\tau} a_k^{p/2}$,

$$\begin{aligned} \mathbb{E} \left[\left| \sum_{t=0}^{\tau_{n+1}} D_{t, \tau_{n+1}}^{N, \theta_n} (\mathbf{S}_{\tau_{n+1}}) \right|^p \right] &\leq C \left(\frac{\sigma_+ |v|_\infty}{\sigma_-^2} \right)^p \\ &\times \frac{(\tau_{n+1} + 1)^{(\frac{p}{2}-1)\vee 0}}{N^p} \sum_{t=0}^{\tau_{n+1}} \mathbb{E} \left[\left| \frac{1}{b_-(\mathbf{Y}_{t+T_n})} \sum_{\ell=1}^N \omega_t^\ell \frac{G_{t, \tau_{n+1}}^{N, \theta_n} \mathbf{S}_{\tau_{n+1}}(\xi_t^\ell)}{|\mathcal{L}_{t, \tau_{n+1}}^{\theta_n} \mathbf{1}|_\infty} \right|^p \right]. \end{aligned}$$

Since $\mathbf{Y}_{t+T_n}$ and θ_n are $\mathcal{F}_{n, t-1}^N$ -measurable,

$$\begin{aligned} \mathbb{E} \left[\left| \frac{1}{b_-(\mathbf{Y}_{t+T_n})} \sum_{\ell=1}^N \omega_t^\ell \frac{G_{t, \tau_{n+1}}^{N, \theta_n} \mathbf{S}_{\tau_{n+1}}(\xi_t^\ell)}{|\mathcal{L}_{t, \tau_{n+1}}^{\theta_n} \mathbf{1}|_\infty} \right|^p \right] \\ = \mathbb{E} \left[\mathbb{E} \left[\left| \sum_{\ell=1}^N \omega_t^\ell \frac{G_{t, \tau_{n+1}}^{N, \theta_n} \mathbf{S}_{\tau_{n+1}}(\xi_t^\ell)}{|\mathcal{L}_{t, \tau_{n+1}}^{\theta_n} \mathbf{1}|_\infty} \right|^p \middle| \mathcal{F}_{n, t-1}^N \right] \frac{1}{b_-(\mathbf{Y}_{t+T_n})^p} \right] \end{aligned}$$

By Lemma B.5(i), using again the Burkholder and convexity inequalities, there exists C s.t.

$$\begin{aligned} \mathbb{E} \left[\left| \sum_{\ell=1}^N \omega_t^\ell \frac{G_{t, \tau_{n+1}}^{N, \theta_n} \mathbf{S}_{\tau_{n+1}}(\xi_t^\ell)}{|\mathcal{L}_{t, \tau_{n+1}}^{\theta_n} \mathbf{1}|_\infty} \right|^p \middle| \mathcal{F}_{t-1, n}^N \right] &\leq C N^{(\frac{p}{2}-1)\vee 0} \mathbb{E} \left[\left| \sum_{\ell=1}^N \omega_t^\ell \frac{G_{t, \tau_{n+1}}^{N, \theta_n} \mathbf{S}_{\tau_{n+1}}(\xi_t^\ell)}{|\mathcal{L}_{t, \tau_{n+1}}^{\theta_n} \mathbf{1}|_\infty} \right|^p \middle| \mathcal{F}_{n, t-1}^N \right] \\ &\leq C N^{\frac{p}{2}\vee 1} \mathbb{E} \left[\left| \omega_t^1 \frac{G_{t, \tau_{n+1}}^{N, \theta_n} \mathbf{S}_{\tau_{n+1}}(\xi_t^1)}{|\mathcal{L}_{t, \tau_{n+1}}^{\theta_n} \mathbf{1}|_\infty} \right|^p \middle| \mathcal{F}_{n, t-1}^N \right]. \end{aligned}$$

The proof is concluded by (43). $\square$

Proposition B.7. *Assume A2 and A6. Let $\{(\xi_t^\ell, \omega_t^\ell), 1 \leq \ell \leq N, 0 \leq t \leq \tau_{n+1}\}$ be the weighted samples obtained by Algorithm 2 in Appendix A, with input variables $\theta_n, \tau_{n+1}, N, \mathbf{Y}_{T_n+1:T_n+\tau_{n+1}}$. For all $\bar{p} > 1$ and all $p \in (1, \bar{p})$, there exists a constant C s.t. for any $t \in \{0, \dots, \tau_{n+1}\}$,*

$$\begin{aligned} \mathbb{E} \left[\left| C_{t, \tau_{n+1}}^{N, \theta_n} (\mathbf{S}_{\tau_{n+1}}) \right|^p \right] &\leq C N^{(\frac{p}{2}\vee \frac{1}{\alpha}) + (\frac{p}{2}\vee \frac{1}{\beta}) - 2p} \\ &\times \mathbb{E} \left[\left| \frac{\omega_+(\mathbf{Y}_{t+T_n})}{b_-(\mathbf{Y}_{t+T_n})} \right|^{2p} \mathbb{E} \left[\left| \sum_{s=1}^{\tau_{n+1}} \rho^{|t-s|} \text{osc}\{S(\cdot, \cdot, \mathbf{Y}_{s+T_n})\} \right|^{\bar{p}} \middle| \mathcal{F}_{n, t-1}^N \right]^{p/\bar{p}} \right], \quad (45) \end{aligned}$$

where $C_{t,\tau}^{N,\theta}$ is defined in (40) and $\alpha \stackrel{\text{def}}{=} \bar{p}/p$ and $\beta^{-1} = 1 - \alpha^{-1}$.

Proof. Lemma B.4 applied with the function $h = \mathcal{L}_{t,\tau_{n+1}}^{\theta_n} \mathbf{1}$ and (36) yield for any $1 \leq \ell \leq N$

$$\mathbb{E} \left[\omega_t^\ell \mathcal{L}_{t,\tau_{n+1}}^{\theta_n} \mathbf{1}(\xi_t^\ell) \middle| \mathcal{F}_{n,t-1}^N \right] = \frac{\phi_{t-1}^{N,\theta_n} \left[\mathcal{L}_{t-1,\tau_{n+1}}^{\theta_n} \mathbf{1} \right]}{\phi_{t-1}^{N,\theta_n} [v_t]}.$$

Therefore, by definition of $C_{t,\tau}^{N,\theta}$ (see (40)), $C_{t,\tau_{n+1}}^{N,\theta_n}(\mathbf{S}_{\tau_{n+1}})$ is equal to

$$\begin{aligned} & \frac{\mathbb{E} [A_{n,t}^N | \mathcal{F}_{n,t-1}^N] - A_{n,t}^N}{\mathbb{E} [A_{n,t}^N | \mathcal{F}_{n,t-1}^N] A_{n,t}^N} B_{n,t}^N = (\mathbb{E} [A_{n,t}^N | \mathcal{F}_{n,t-1}^N] - A_{n,t}^N) \cdots \\ & \times \frac{\Omega_{n,t}^N}{\mathbb{E} [A_{n,t}^N | \mathcal{F}_{n,t-1}^N] A_{n,t}^N} \left(\frac{B_{n,t}^N}{\mathbb{E} [\Omega_{n,t}^N | \mathcal{F}_{n,t-1}^N]} + \frac{B_{n,t}^N}{\Omega_{n,t}^N \mathbb{E} [\Omega_{n,t}^N | \mathcal{F}_{n,t-1}^N]} (\mathbb{E} [\Omega_{n,t}^N | \mathcal{F}_{n,t-1}^N] - \Omega_{n,t}^N) \right) \\ & = B_{n,t}^N (\mathbb{E} [A_{n,t}^N | \mathcal{F}_{n,t-1}^N] - A_{n,t}^N) \frac{\Omega_{n,t}^N}{\mathbb{E} [A_{n,t}^N | \mathcal{F}_{n,t-1}^N] A_{n,t}^N} \frac{1}{\mathbb{E} [\Omega_{n,t}^N | \mathcal{F}_{n,t-1}^N]} \\ & + (\mathbb{E} [A_{n,t}^N | \mathcal{F}_{n,t-1}^N] - A_{n,t}^N) (\mathbb{E} [\Omega_{n,t}^N | \mathcal{F}_{n,t-1}^N] - \Omega_{n,t}^N) \frac{B_{n,t}^N}{\mathbb{E} [A_{n,t}^N | \mathcal{F}_{n,t-1}^N] A_{n,t}^N} \frac{1}{\mathbb{E} [\Omega_{n,t}^N | \mathcal{F}_{n,t-1}^N]} \end{aligned}$$

with

$$\begin{aligned} A_{n,t}^N & \stackrel{\text{def}}{=} N^{-1} \sum_{\ell=1}^N \omega_t^\ell \frac{\mathcal{L}_{t,\tau_{n+1}}^{\theta_n} \mathbf{1}(\xi_t^\ell)}{|\mathcal{L}_{t,\tau_{n+1}}^{\theta_n} \mathbf{1}|_\infty}, \\ B_{n,t}^N & \stackrel{\text{def}}{=} \frac{1}{N} \sum_{\ell=1}^N \omega_t^\ell \frac{G_{t,\tau_{n+1}}^{N,\theta_n} \mathbf{S}_{\tau_{n+1}}(\xi_t^\ell)}{|\mathcal{L}_{t,\tau_{n+1}}^{\theta_n} \mathbf{1}|_\infty}, \\ \Omega_{n,t}^N & \stackrel{\text{def}}{=} \frac{1}{N} \sum_{\ell=1}^N \omega_t^\ell. \end{aligned}$$

This can be rewritten,

$$C_{t,\tau}^{N,\theta} = C_1 + C_2,$$

with

$$C_1 = B_{n,t}^N (\mathbb{E} [A_{n,t}^N | \mathcal{F}_{n,t-1}^N] - A_{n,t}^N) \frac{\Omega_{n,t}^N}{\mathbb{E} [A_{n,t}^N | \mathcal{F}_{n,t-1}^N] A_{n,t}^N} \frac{1}{\mathbb{E} [\Omega_{n,t}^N | \mathcal{F}_{n,t-1}^N]}$$

and

$$\begin{aligned} C_2 & = (\mathbb{E} [A_{n,t}^N | \mathcal{F}_{n,t-1}^N] - A_{n,t}^N) (\mathbb{E} [\Omega_{n,t}^N | \mathcal{F}_{n,t-1}^N] - \Omega_{n,t}^N) \\ & \times \frac{B_{n,t}^N}{\mathbb{E} [A_{n,t}^N | \mathcal{F}_{n,t-1}^N] A_{n,t}^N} \frac{1}{\mathbb{E} [\Omega_{n,t}^N | \mathcal{F}_{n,t-1}^N]}. \end{aligned}$$

By Lemmas B.4 and Lemmas B.5(iii), and A6,

$$\frac{1}{\mathbb{E} [\Omega_{n,t}^N | \mathcal{F}_{n,t-1}^N]} \leq \frac{\sigma_- |v|_\infty}{b_- (\mathbf{Y}_{t+T_n})}; \quad \mathbb{E} \left[\frac{\Omega_{t,n}^N}{A_{n,t}^N | \mathcal{F}_{n,t-1}^N} \right] \leq \left(\frac{\sigma_+}{\sigma_-} \right)^2 \frac{|v|_\infty}{\sigma_- b_- (\mathbf{Y}_{t+T_n})};$$

and by Lemma B.5(ii)

$$\frac{B_{n,t}^N}{\mathbb{E} [A_{n,t}^N | \mathcal{F}_{n,t-1}^N] A_{n,t}^N} \leq \left(\frac{\sigma_+}{\sigma_-} \right)^2 \frac{|v|_\infty}{\sigma_- b_- (\mathbf{Y}_{t+T_n})} \left(\sum_{s=1}^{\tau_{n+1}} \rho^{|t-s|} \text{osc}(S(\cdot, \cdot, \mathbf{Y}_{s+T_n})) \right).$$

Therefore, there exists a constant C s.t.

$$\mathbb{E} [|C_1|^p | \mathcal{F}_{n,t-1}^N] \leq C \left| \frac{1}{b_- (\mathbf{Y}_{t+T_n})} \right|^{2p} \mathbb{E} \left[|B_{n,t}^N|^p | \mathbb{E} [A_{n,t}^N | \mathcal{F}_{n,t-1}^N] - A_{n,t}^N |^p | \mathcal{F}_{n,t-1}^N \right]$$

Applying the Holder inequality with $\alpha \stackrel{\text{def}}{=} \bar{p}/p \geq 1$ and $\beta^{-1} \stackrel{\text{def}}{=} 1 - \alpha^{-1}$ yields

$$\begin{aligned} & \mathbb{E} \left[|B_{n,t}^N|^p | \mathbb{E} [A_{n,t}^N | \mathcal{F}_{n,t-1}^N] - A_{n,t}^N |^p | \mathcal{F}_{n,t-1}^N \right] \\ & \leq \mathbb{E} \left[|B_{n,t}^N|^{\alpha p} | \mathcal{F}_{n,t-1}^N \right]^{1/\alpha} \mathbb{E} \left[|\mathbb{E} [A_{n,t}^N | \mathcal{F}_{n,t-1}^N] - A_{n,t}^N|^{\beta p} | \mathcal{F}_{n,t-1}^N \right]^{1/\beta}. \end{aligned}$$

By Proposition B.6,

$$\begin{aligned} & \mathbb{E} \left[|B_{n,t}^N|^{\alpha p} | \mathcal{F}_{n,t-1}^N \right]^{1/\alpha} \\ & \leq C N^{(\frac{p}{2} \vee \frac{1}{\alpha}) - p} \mathbb{E} \left[\left| \omega_t^1 \frac{G_{t,\tau_{n+1}}^{N,\theta_n} S_{\tau_{n+1}}(\xi_t^1)}{|\mathcal{L}_{t,\tau_{n+1}}^{\theta_n} \mathbf{1}|_\infty} \right|^{\alpha p} | \mathcal{F}_{n,t-1}^N \right]^{1/\alpha} \\ & \leq C N^{(\frac{p}{2} \vee \frac{1}{\alpha}) - p} \omega_+ (\mathbf{Y}_{t+T_n})^p \mathbb{E} \left[\left| \sum_{s=1}^{\tau_{n+1}} \rho^{|t-s|} \text{osc}\{S(\cdot, \cdot, \mathbf{Y}_{s+T_n})\} \right|^{\alpha p} | \mathcal{F}_{n,t-1}^N \right]^{1/\alpha}, \end{aligned}$$

Given $\mathcal{F}_{n,t-1}^N$, the random variables $\left\{ \mathbb{E} \left[\omega_t^\ell \frac{\mathcal{L}_{t,\tau_{n+1}}^{\theta_n} \mathbf{1}(\xi_t^\ell)}{|\mathcal{L}_{t,\tau_{n+1}}^{\theta_n} \mathbf{1}|_\infty} | \mathcal{F}_{n,t-1}^N \right] - \omega_t^\ell \frac{\mathcal{L}_{t,\tau_{n+1}}^{\theta_n} \mathbf{1}(\xi_t^\ell)}{|\mathcal{L}_{t,\tau_{n+1}}^{\theta_n} \mathbf{1}|_\infty} \right\}_{\ell=1}^N$ are conditionally independent, centered and bounded by Lemma B.5. Following the same steps as in the proof of Proposition B.6, there exists a constant C such that

$$\mathbb{E} \left[|\mathbb{E} [A_{n,t}^N | \mathcal{F}_{n,t-1}^N] - A_{n,t}^N|^{\beta p} | \mathcal{F}_{n,t-1}^N \right]^{1/\beta} \leq C N^{(\frac{p}{2} \vee \frac{1}{\beta}) - p} \omega_+ (\mathbf{Y}_{t+T_n})^p.$$

Hence,

$$\begin{aligned} & \mathbb{E} [|C_1|^p] \leq C N^{(\frac{p}{2} \vee \frac{1}{\alpha}) + (\frac{p}{2} \vee \frac{1}{\beta}) - 2p} \\ & \times \mathbb{E} \left[\left| \frac{\omega_+ (\mathbf{Y}_{t+T_n})}{b_- (\mathbf{Y}_{t+T_n})} \right|^{2p} \mathbb{E} \left[\left| \sum_{s=1}^{\tau_{n+1}} \rho^{|t-s|} \text{osc}\{S(\cdot, \cdot, \mathbf{Y}_{s+T_n})\} \right|^{\bar{p}} | \mathcal{F}_{n,t-1}^N \right]^{p/\bar{p}} \right]. \end{aligned}$$

Similarly, using

$$\mathbb{E} \left[\left| \mathbb{E} [\Omega_{t,n}^N | \mathcal{F}_{n,t-1}^N] - \Omega_{t,n}^N \right|^{\alpha p} \middle| \mathcal{F}_{n,t-1}^N \right]^{1/\alpha} \leq C N^{(\frac{p}{2} \vee \frac{1}{\alpha}) - p} \omega_+(\mathbf{Y}_{t+T_n})^p ,$$

yields

$$\begin{aligned} \mathbb{E} [|C_2|^p] &\leq C N^{-p} \\ &\times \mathbb{E} \left[\left| \frac{\omega_+(\mathbf{Y}_{t+T_n})}{b_-(\mathbf{Y}_{t+T_n})} \right|^{2p} \mathbb{E} \left[\left| \sum_{s=1}^{\tau_{n+1}} \rho^{|t-s|} \text{osc}\{S(\cdot, \cdot, \mathbf{Y}_{s+T_n})\} \right|^{\bar{p}} \middle| \mathcal{F}_{n,t-1}^N \right]^{p/\bar{p}} \right] . \end{aligned}$$

□

References

- [1] T. Bailey, J. Nieto, J. Guivant, M. Stevens, and E. Nebot. Consistency of the EKF-SLAM algorithm. In *IEEE International Conference on Intelligent Robots and Systems*, pages 3562–3568, 2006.
- [2] W. Burgard, D. Fox, and S. Thrun. *Probabilistic robotics*. Cambridge, MA:MIT Press, 2005.
- [3] O. Cappé. Online sequential Monte Carlo EM algorithm. In *IEEE Workshop on Statistical Signal Processing (SSP)*, pages 37–40, 2009.
- [4] O. Cappé. Online EM algorithm for Hidden Markov Models. *To appear in J. Comput. Graph. Statist.*, 2011.
- [5] O. Cappé and E. Moulines. Online Expectation Maximization algorithm for latent data models. *J. Roy. Statist. Soc. B*, 71(3):593–613, 2009.
- [6] O. Cappé, E. Moulines, and T. Rydén. *Inference in Hidden Markov Models*. Springer, New York, 2005.
- [7] J. Davidson. *Stochastic Limit Theory: An Introduction for Econometricians*. Oxford University Press, 1994.
- [8] P. Del Moral. *Feynman-Kac Formulae. Genealogical and Interacting Particle Systems with Applications*. Springer, New York, 2004.
- [9] P. Del Moral, A. Doucet, and S.S. Singh. A Backward Particle Interpretation of Feynman-Kac Formulae. *ESAIM M2AN*, 44(5):947–975, 2010.
- [10] P. Del Moral, A. Doucet, and S.S. Singh. Forward smoothing using sequential Monte Carlo. Technical report, arXiv:1012.5390v1, 2010.
- [11] P. Del Moral and A. Guionnet. Large deviations for interacting particle systems: applications to non-linear filtering. *Stoch. Proc. App.*, 78:69–95, 1998.

- [12] P. Del Moral, M. Ledoux, and L. Miclo. On contraction properties of Markov kernels. *Probab. Theory Related Fields*, 126(3):395–420, 2003.
- [13] B. Delyon, M. Lavielle, and E. Moulines. Convergence of a stochastic approximation version of the EM algorithm. *Ann. Statist.*, 27(1), 1999.
- [14] A. P. Dempster, N. M. Laird, and D. B. Rubin. Maximum likelihood from incomplete data via the EM algorithm. *J. Roy. Statist. Soc. B*, 39(1):1–38 (with discussion), 1977.
- [15] R. Douc, G. Fort, E. Moulines, and P. Priouret. Forgetting the initial distribution for hidden Markov models. *Stochastic Processes and their Applications*, 119(4):1235–1256, 2009.
- [16] R. Douc, A. Garivier, E. Moulines, and J. Olsson. Sequential Monte Carlo smoothing for general state space hidden Markov models. *To appear in Ann. Appl. Probab.*, 4 2010.
- [17] A. Doucet, N. De Freitas, and N. Gordon, editors. *Sequential Monte Carlo Methods in Practice*. Springer, New York, 2001.
- [18] A. Doucet, G. Poyiadjis, and S.S. Singh. Particle approximations of the score and observed information matrix in state-space models with application to parameter estimation. *Biometrika*, 98(1):65–80, 2010.
- [19] C. Dubarry and S. Le Corff. Non-asymptotic deviation inequalities for smoothed additive functionals in non-linear state-space models. Technical report, arXiv:1012.4183v1, 2011.
- [20] G. Fort and E. Moulines. Convergence of the Monte Carlo Expectation Maximization for curved exponential families. *Ann. Statist.*, 31(4):1220–1259, 2003.
- [21] P. Hall and C. C. Heyde. *Martingale Limit Theory and its Application*. Academic Press, New York, London, 1980.
- [22] S. Le Corff and G. Fort. Convergence of a particle-based approximation of the Block Online Expectation Maximization algorithm. Technical report, arXiv, 2011.
- [23] S. Le Corff and G. Fort. Online Expectation Maximization based algorithms for inference in Hidden Markov Models. Technical report, arXiv:1108.3968v1, 2011.
- [24] S. Le Corff and G. Fort. Supplement paper to "Online Expectation Maximization based algorithms for inference in Hidden Markov Models". Technical report, arXiv:1108.4130v1, 2011.
- [25] S. Le Corff, G. Fort, and E. Moulines. Online EM algorithm to solve the SLAM problem. In *IEEE Workshop on Statistical Signal Processing (SSP)*, 2011.

- [26] F. Le Gland and L. Mevel. Recursive estimation in HMMs. In *Proc. IEEE Conf. Decis. Control*, pages 3468–3473, 1997.
- [27] J.S. Liu. *Monte Carlo Strategies in Scientific Computing*. Springer, New York, 2001.
- [28] Ruben Martinez-Cantin. *Active Map Learning for Robots: Insights into Statistical Consistency*. PhD thesis, University of Zaragoza, 2008.
- [29] G. J. McLachlan and T. Krishnan. *The EM Algorithm and Extensions*. Wiley, New York, 1997.
- [30] G. J. McLachlan and S. K. NG. On the choice of the number of blocks with the incremental EM algorithm for the fitting of normal mixtures. *Statistics and Computing*, 13(1):45–55, 2003.
- [31] S. P. Meyn and R. L. Tweedie. *Markov Chains and Stochastic Stability*. Springer, London, 1993.
- [32] M. Montemerlo, S. Thrun, D. Koller, and B. Wegbreit. FastSLAM 2.0: An improved particle filtering algorithm for simultaneous localization and mapping that provably converges. In *Proceedings of the Sixteenth IJCAI*, Mexico, 2003.
- [33] B. T. Polyak. A new method of stochastic approximation type. *Autom. Remote Control*, 51:98–107, 1990.
- [34] Vladislav B. Tadić. Analyticity, convergence, and convergence rate of recursive maximum-likelihood estimation in hidden Markov models. *IEEE Trans. Inf. Theor.*, 56:6406–6432, December 2010.
- [35] D. M. Titterton. Recursive parameter estimation using incomplete data. *J. Roy. Statist. Soc. B*, 46(2):257–267, 1984.
- [36] C. F. J. Wu. On the convergence properties of the EM algorithm. *Ann. Statist.*, 11:95–103, 1983.